

Decoupled Method for Reconstruction of Surface Conditions From Internal Temperatures On Ablative Materials With Uncertain Recession Model

A. Brandon Oliver,*

NASA Lyndon B. Johnson Space Center

Houston, TX, 77058

Obtaining measurements of flight environments on ablative heatshields is both critical for spacecraft development and extremely challenging due to the harsh heating environment and surface recession. Thermocouples installed several millimeters below the surface are commonly used to measure the heatshield temperature response, but an ill-posed inverse heat conduction problem must be solved to reconstruct the surface heating environment from these measurements. Ablation can contribute substantially to the measurement response making solutions to the inverse problem strongly dependent on the recession model, which is often poorly characterized. To enable efficient surface reconstruction for recession model sensitivity analysis, a method for decoupling the surface recession evaluation from the inverse heat conduction problem is presented. The decoupled method is shown to provide reconstructions of equivalent accuracy to the traditional coupled method but with substantially reduced computational effort. These methods are applied to reconstruct the environments on the Mars Science Laboratory heatshield using diffusion limit and kinetically limited recession models.

I. Introduction

I.A. Motivation

Partial validation of aerothermal environments developed for atmospheric entry of a spacecraft can be performed in carefully scaled ground tests; however, only flight can capture all of the relevant physics and their interactions, making flight data extremely valuable to the design and development of a spacecraft. When flight data is used to validate environments used to size a thermal protection system (TPS) heatshield, measurement and processing errors can contribute to the likelihood of heatshield failure or over-conservatism if they falsely ‘validate’ bad predictions or ‘invalidate’ good predictions. Care must be taken to reduce and quantify experimental errors to every extent possible.

Heat flux can be a difficult quantity to measure as it can only be inferred from measurements of other properties. All sensors that presently ‘measure’ heat flux operate by measuring temperature at one or more locations and then inferring the heat flux from the temperatures and thermal response assumptions (either explicitly modeled in a math-model or empirically modeled through calibration). This reliance on inference places restrictions on the operational conditions of heat flux sensors. There are few sensors that can operate accurately at the conditions seen in atmospheric reentry.

The windward side heatshields for NASA’s Orion and Mars Science Laboratory (MSL) capsules use charring ablaters, so it is expected that the heatshield will ablate and recede, leaving anything embedded in the heatshield to protrude into the oncoming flow. A protruding sensor can amplify heating in the vicinity of the protrusion and spoil the measurement. The Apollo Program instrumented a few of the flight test vehicle ablative heatshields with sacrificial calorimeters¹, but this path was not followed for the Orion EFT-1 flight test, nor the MSL Entry Descent and Landing Instrumentation (MEDLI) program²⁻⁵. These vehicles used embedded thermocouples (TCs) inside their respective heatshields with the TCs set deep enough that they would survive through the relevant part of reentry. As with the other types of heat flux sensor, the surface heating must be inferred from the actual sensor measurements, in this case by solving a problem known as the inverse heat conduction problem (IHCP). The solution of the inverse problem, a process which is often referred to as *surface condition reconstruction*, is not always straightforward, especially on ablaters.

*Applied Aeroscience and CFD Branch.

I.B. Reconstruction of Surface Conditions on Ablators

The IHCP⁶⁻⁸, or *inverse* problem, can be stated mathematically as:

Given:

$$C \frac{\partial T}{\partial t} = \frac{\partial}{\partial x} \left(k \frac{\partial T}{\partial x} \right) \quad (1a)$$

$$-k \frac{\partial T}{\partial x} \Big|_{x=x_s} = q(t) \quad (1b)$$

$$T(L, t) = g(t) \quad (1c)$$

$$T(x, 0) = f(x) \quad (1d)$$

$$T(x_i, t_j) = y_{ij} + \epsilon_{ij} \quad (1e)$$

Find:

$$q(t) \text{ for } t > 0 \quad (2)$$

where coefficients C , k , L , x_s , and T_∞ are assumed known, as are the measurement locations x_j , times t_i , and values y_{ij} (to some limited accuracy because of unknown ϵ_{ij}), functions $g(t)$ and $f(x)$ are assumed known, but $q(t)$ is unknown (note that different boundary conditions than specified above can be used; these are included as common examples). In a *direct* or *forward* problem, $q(t)$ is known and $T(x, t)$ is sought. The inverse problem is ill-posed, meaning that the solution is not guaranteed to be unique and does not always vary smoothly with small perturbations on the inputs. Optimization methods⁷⁻¹¹ are frequently used to obtain a solution for the unknown boundary condition that minimizes the least-squares difference between measured temperatures and temperatures calculated from a material response model. Any errors in the material response model will affect the reconstructed boundary condition.

On an ablator, the boundary location x_s varies and must also be determined. The evolution of the surface location is governed by a *recession model*. The recession model may define the ablation rate \dot{m}_c'' with an empirical relationship based on surface conditions or an explicit chemical reaction mechanism, but it often defines the ablation rate as that which yields heterogeneous and homogeneous chemical equilibrium at the surface of the ablator (called a *diffusion limit* model). The physical mechanisms of carbon ablation is an area of active research, and most models being proposed require tight coupling between CFD and ablation response. At the present time, this capability is not developed enough to be a feasible means of analysis on a problem of the scale and complexity of a full heatshield. Instead, the Orion and MSL heatshield ablation analysis is performed using Apollo-era models of ablation¹²⁻¹⁷. These models introduce a number of assumptions¹⁸ to decouple the flowfield, ablation, and thermal response models. The decoupling relies on the film coefficient model given by

$$C_h^* = \frac{\dot{q}_{aero}''}{H_{rec} - H_w}, \quad (3)$$

where \dot{q}_{aero}'' is the net heat flux from the boundary layer, C_h^* is the film coefficient, H_{rec} is the boundary layer recovery enthalpy, and H_w is the wall enthalpy (enthalpy of ablation products), to scale the heat and mass transfer between the flowfield and surface. The film coefficient model appears in the engineering-form of the *surface energy balance* (SEB)

$$\dot{q}_{cond_s}'' = \overbrace{C_h^*(H_{rec} - H_w)}^{\text{Environment}} + \overbrace{\alpha \dot{q}_{rad}'' - \sigma \epsilon (T_s^4 - T_\infty^4)}^{\text{Reradiation}} + \overbrace{\dot{m}_c'' H_c + \dot{m}_g'' H_g - (\dot{m}_c'' + \dot{m}_g'') H_w - \dot{m}_f'' H_f}_{\text{Ablation Terms}}, \quad (4)$$

with \dot{q}_{cond_s}'' representing the heat flux into the surface through conduction, α the surface absorptivity, \dot{q}_{rad}'' the shock-layer radiation, σ the Boltzman constant, ϵ the surface emissivity, T_s and T_∞ the surface and far-field temperatures respectively, \dot{m}_c'' the ablation rate, H_c the enthalpy of the char, \dot{m}_g'' the pyrolysis gas blowing rate, H_g the pyrolysis gas enthalpy, \dot{m}_f'' the fail-ablation rate, and H_f the fail-material enthalpy. The terms highlighted in blue are typically provided by separate CFD or boundary layer analysis, and the red terms are defined by the recession model. The remaining right-hand-side terms are material properties (assumed known, possibly functions of temperature and pressure) or evaluated by the material response model. Evaluation of the recession model at the appropriate surface conditions yields the red-highlighted terms needed to complete the SEB and evaluate the conduction heat flux boundary condition required for the thermal response analysis. The recession model also provides the ablation rate needed to update the surface location.

For recession models that assume a chemical reaction mechanism or chemical equilibrium at the surface, the film coefficient is one of the required surface conditions since it characterizes the flux of boundary layer species to the surface. As a result, solutions to the inverse problem must be formulated to solve for the film coefficient on a domain that evolves based on the evaluated film coefficient (i.e. *find* $C_h^*(t)$ when $q(t)$ and x_s in Equation 1b are given by $q(t) = \dot{q}_{cond_s}(C_h^*)$ and $x_s(C_h^*)$). This can pose a number of challenges to an IHCP algorithm.

First of all, an ablating surface can recess past an embedded thermocouple (referred to as burn-out) either in the physical experiment or in the material response model. This is not necessarily a problem if the recession model is physically accurate. However, the recession model is often poorly characterized and introduces non-negligible modeling errors that lead to a difference in burn-out times. Once burn-out occurs, comparisons between measurement and model cannot be made and the reconstruction will be unable to continue. Secondly, variations in the various terms of the SEB can lead the TC response to be more or less sensitive to the film coefficient at different times in the reconstruction, the extent of which may be dependent on the film coefficient. This increases the non-linearity of the IHCP. In extreme cases, such as those experienced on MSL^{19,20}, the TCs can become almost completely insensitive to the film coefficient. Note that the film coefficient is multiplied by the term $(H_{rec} - H_w)$ in Equation 4; if this term goes to zero, the TCs will become insensitive to C_h^* and the reconstruction will likely become unstable. Finally, inaccurate estimation of the enthalpy of the ablation products can drive errors into the reconstructed film coefficient in order to produce the appropriate heat flux required from the first term of the SEB.

I.C. Characterization of Modeling Assumptions on Reconstructed Film Coefficient

If a reconstruction of flight measurements is to be used to validate predicted heating environments, an uncertainty analysis of the reconstruction should be performed. Uncertainty propagation methods like those of Blackwell et al.²¹ are useful for linear and nearly-linear problems; however, the material response of typical charring ablators are too non-linear for methods such as this to provide sufficient characterization of the reconstruction uncertainty. Sampling methods (such as the Monte Carlo or Latin Hypercube Sampling techniques) with subsequent statistical analysis provide the most straightforward method of assessing uncertainties in material response model terms. While relatively simple to implement, these methods can be quite computationally expensive, as hundreds if not thousands of reconstructions must be performed to adequately address all of the material response model inputs. Since the cost of a single IHCP solution on an ablator can consume several hundred CPU-hours, the cost of reconstruction must be reduced to permit practical evaluation of uncertainty through sampling methods.

I.D. Objective

To address some of the limitations mentioned above, a method is proposed in this paper whereby the surface energy balance solution (and any associated surface recession) is decoupled from the IHCP solution and solved in a separate step following the IHCP reconstruction. The theoretical basis for this method is presented, the code implementation is briefly described, and two verification examples and a comparison to a conventional coupled reconstruction are presented to demonstrate the accuracy of the method. A subset of the MEDLI flight data will be reconstructed to demonstrate and highlight the efficiency of the method on a real problem.

II. Decoupling Theory

Consider the non-linear 1-D heat equation

$$C \frac{\partial T}{\partial t} - \frac{\partial}{\partial x} \left(k \frac{\partial T}{\partial x} \right) = 0 \quad (5a)$$

on the domain

$$x_0 \leq x \leq L \quad (5b)$$

with boundary and initial conditions given by

$$-k \frac{\partial T}{\partial x} \Big|_{x_0} = \dot{q}''(t) \quad (5c)$$

$$T(L, t) = g(t) \quad (5d)$$

$$T(x, 0) = f(x) \quad (5e)$$

that satisfy all conditions necessary to yield a unique solution for $T(x, t)$ for $t > 0$. Two time-varying locations on the domain are defined such that

$$x_0 \leq x_s(t) \leq x_m(t) \leq L, \quad (6)$$

for $t > 0$ (the temporal dependance is omitted from subsequent nomenclature for clarity). If a second similar problem is defined as

$$C \frac{\partial Y}{\partial t} - \frac{\partial}{\partial x} \left(k \frac{\partial Y}{\partial x} \right) = 0 \quad (7a)$$

on the restricted spatial domain

$$x_s \leq x \leq L \quad (7b)$$

with boundary and initial conditions given by

$$Y(x_s, t) = T(x_s, t) \quad (7c)$$

$$Y(L, t) = g(t) \quad (7d)$$

$$Y(x, 0) = f(x), \quad (7e)$$

and C , k , g , and f identical to those in the first problem (only Equation 7c changes), then the uniqueness of the solution to Equation 5 implies that

$$Y(x_m, t) = T(x_m, t) \quad (8)$$

for any x_m on $[x_s, L]$ for $t > 0$. In the subsequent discussion, the system given by Equation 5 is referred to as the *full system*, and the system given by Equation 7 is referred to as the *restricted system*.

Taken one step further, if the two problems are identical except for the extent of the spatial domain and if the equality of Equation 8 is enforced, then the uniqueness properties of the heat equation solution require that the boundary condition on the surface of the restricted system must be given by Equation 7c (or the Neumann equivalent). Notice that stated this way, the problem takes the general form of an inverse problem. Also notice that the only restrictions on the interior points x_s and x_m above are that they are on the domain and that $x_s \leq x_m$. The implication of this is that an IHCP reconstruction is not only a reconstruction of the boundary condition, but a reconstruction of the whole temperature field *consistent with the governing equation and measurements at a point, x_m , on the domain*.

This result can be used in the reconstruction of surface conditions when the true surface location is unknown at the time of reconstruction. The restricted system represents the true physical system with the surface at x_s and the measurements describing $Y(x_m, t)$. IHCP algorithms are used to reconstruct an appropriate boundary condition at x_0 ($\dot{q}''(t)$ in Equation 5c) for the full system that yields $T(x_m, t) = Y(x_m, t)$. The resulting solution of the full system can then provide the T and $\frac{\partial T}{\partial x}$ needed to compute \dot{q}''_{cond_s} in the surface energy balance (Equation 4) for any possible surface location on $x_0 \leq x_s \leq x_m$ that is consistent with the measurements. Any information regarding the location of x_s can then be incorporated to define the reconstructed surface conditions, or a separate model describing the motion of x_s with time (even one that is a function of T and $\frac{\partial T}{\partial x}$ at x_s) can be evaluated separate from the IHCP solution. The key is that the temperature field reconstruction is decoupled from the surface recession evaluation.

These concepts are illustrated graphically in Figures 1 and 2. Figure 1 shows the temperature solution at time t_1 of the full system that has been reconstructed to provide the measured temperature at the measurement location x_m . Two potential values of x_s are shown as x_1 and x_2 . Under a set of assumptions that specifies x_1 as the true surface location, the surface temperature must be 1200 K in order to provide the proper response at the sensor. In this case, the temperature at x_2 is simply an internal temperature. Alternatively, under a different set of assumptions that suggest x_2 is the true surface location, the surface temperature must be 900 K to satisfy the governing equation and match the sensor measurement. In this case, the temperature at x_1 is outside the true domain and is un-physical under these recession assumptions. Figure 2 shows a plot of the entire temperature field, with contours of both full (thin black lines) and restricted (colored lines) systems co-plotted. The thick black line indicates the location of the true surface defining the restricted system. A slice through this surface at constant depth (magenta line) shows the temperature profile (i.e. the embedded TC measurement) that, along with the governing equations, defines the field.

These two figures highlight a significant caveat that should be addressed: the solutions on the extended domain ($x_0 \leq x < x_s$) are fictitious with respect to the real system and they can include physically impossible temperatures (for example on a melting material, temperatures in this region will exceed the melt temperature). The uniqueness

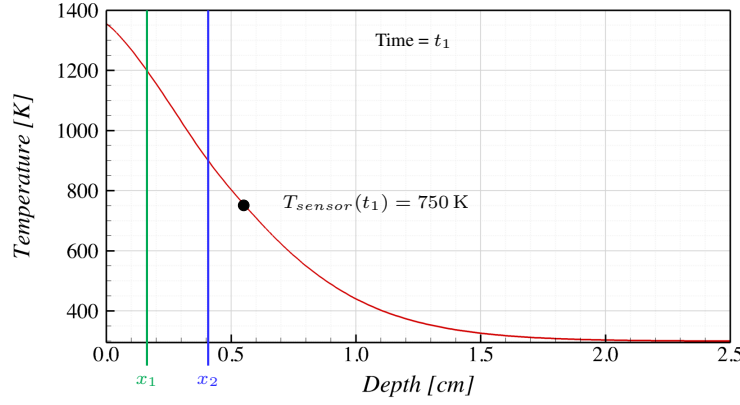


Figure 1: Illustration of a heat equation solution at a specific time showing how different combinations of boundary location and boundary value could yield an internal temperature value consistent with measurement.

property only requires that the restricted system solution and any full system solution be equal on the restricted domain. If C and k are functions of temperature and the full system solution on the extended domain includes values of T that are outside the range of Y , then it may be possible to obtain multiple different solutions to the full system that differ only on the extended domain if different functions for $C(T)$ and $k(T)$ are used (clearly, these functions will only differ for values of T outside the range of Y). Since material properties at non-physical temperatures are non-physical themselves, they may be carefully chosen to improve the stability of the full system inverse reconstruction.

II.A. Pyrolysis Gas Decoupling

It has been shown that the temperature field reconstruction can be computed without knowledge of the exact surface location. Decomposing ablators will have an additional partial differential equation governing the flow of pyrolysis gas. In order to fully reconstruct the behavior of these materials, additional constraints must be addressed. Development of models for the true physical nature of the pyrolysis gas flow is an area of open research; however, there are two common modeling assumptions made regarding the flow of pyrolysis gas relevant here: zero residence time and steady porous flow.

Older ablation models used for the majority of engineering-level ablation analysis (such as CMA¹² and FIAT¹⁷) assume that the pyrolysis gas residence time in the interior of the ablator domain is negligible. The decomposition models are integrated through the domain and any pyrolysis gas produced is assumed to be immediately present at the surface. With this assumption, there is no additional partial differential equation required, and a reconstruction can be completed as previously discussed.

As modeling capability is developing to handle multi-dimensional problems, this assumption becomes difficult to implement since it is not clear where on the surface the pyrolysis gas should emerge from. Consequently, governing equations for the pyrolysis gas flow are necessary. The *CHAR* code^{15,16} used here assumes that the pyrolysis gas adheres to the steady Darcy's law for flow in porous media which relates the gas flow rate to the gradient of the pressure in the ablator pore space (the *internal pressure gradient*) according to

$$\dot{m}_g'' = -\rho_g \frac{\tilde{\kappa}}{\mu} \nabla P, \quad (9)$$

where \dot{m}_g'' is the local mass flux of pyrolysis gas, ρ_g is the gas density, μ is the gas viscosity, $\tilde{\kappa}$ is the permeability tensor of the porous medium, and P is the internal pressure. While the details of derivation are left to reference²², *CHAR* solves an additional parabolic partial differential equation for the internal pressure field that satisfies the gas continuity (conservation of mass) equation. With this approach, the pyrolysis gas can stay resident in the ablator for a finite amount of time, and the flow rate (and direction in multi-dimensional problems) is dependent on permeability. The permeability is generally defined to be a function of the state of decomposition, with higher permeability in decomposed char than in virgin material. When *CHAR* is used to calculate the response of a material for which the model was developed with the zero residence time assumption, the permeability of the virgin and char states are set as low and as high as possible, respectively, that provide stable solutions. In this way, the pyrolysis gas residence time is minimized.

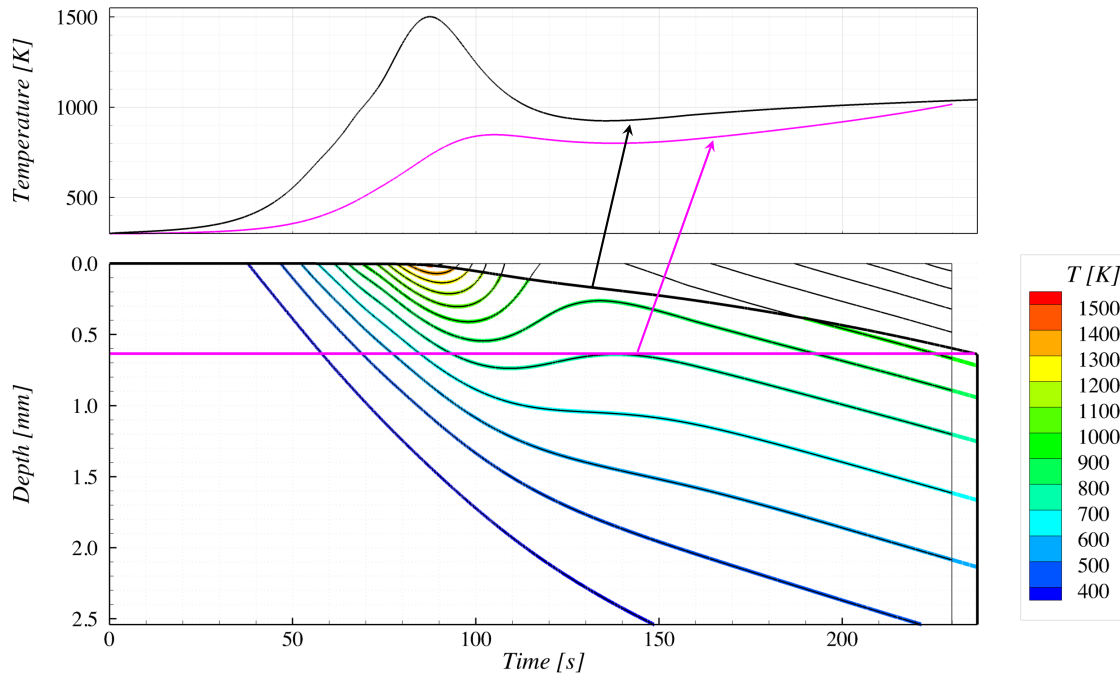


Figure 2: Illustration of a temperature field with a recessed surface shown relative to the fictive non-recessed surface. Note that a cut at a constant depth gives a temperature trace that corresponds to the expected response of a TC at that depth (shown as magenta lines), while the thick black line indicates the location of and conditions on the recessing surface.

For the decoupled SEB approach proposed here, the pyrolysis gas mass flow rate at the surface must be determined since that will likely influence any thermochemical ablation model. For a model that implements additional governing equations, the pyrolysis mass flux must be determined from the coupled solution of the governing equations. This introduces a scenario in which the decoupling approach may break down.

Depending on the material permeability, the internal pressure gradient can be large and non-linear. By itself, this does not invalidate the decoupling theory, at least for the steady porous flow assumption, as the pressure field is governed by a parabolic PDE similar to the temperature field. However, in most reconstructions, the pressure will not be known at an internal location like the temperature will, the pressure will generally only be known on the true surface. Using the known pressure boundary condition on the full system boundary with a material having low permeability will yield internal pressures that are too high relative to the case where the known pressure boundary condition is applied on the boundary of the restricted system. If the pyrolysis gas enthalpy or the solid thermal conductivity are functions of pressure, errors will be introduced (although they will likely be small as these properties generally vary with the log of pressure). More significantly for this case, though, is that if the mass flux field is not uniform in the vicinity of x_s , then the decoupled reconstruction may not be able to define \dot{m}''_g with sufficient accuracy.

For the MEDLI data being addressed in this work, the PICA material model was developed in FIAT using the zero residence time assumption. As such, the *CHAR* implementation of the model uses an artificially high permeability in the char, which leads to a low and linear pressure gradient that minimizes the effect of this modeling assumption. However, a different approach may be required to perform this sort of analysis with future material models that use more realistic permeability models.

II.B. Previous Uses of Decoupled SEB Reconstruction

This approach of decoupling the surface recession calculation from the IHCP solution has been used to some extent in two previous studies of reconstruction of environments on ablative heatshield materials. These are presently described and the differences proposed in this study are outlined.

Mahzari et al.^{19,23} encountered problems with the recession model in their initial attempts to reconstruct the film coefficient from the MEDLI flight data. In collaboration with this group, it was proposed that reconstructing on

a system with suppressed recession would be a meaningful bounding case, and the proposed decoupling approach above was identified. Mahzari et al.²⁰ used this idea to show bounding environments and assessed the sensitivity of a combined convection and ablation “heat rate” given by $\dot{q}_s'' = C_h^*(H_{rec} - H_w) + \dot{m}_c'' H_c + \dot{m}_g'' H_g - (\dot{m}_c'' + \dot{m}_g'') H_w$ for several assumed recession profiles using the temperature derivative evaluated from the no-recession reconstruction. The assumed recession profiles were defined by using uniformly-scaled pre-flight recession estimates. As they did not evaluate a recession model needed to calculate the ablation terms, they could not isolate the film coefficient in their decoupled reconstructions.

Frankel et al.²⁴ have recently proposed a similar two-step reconstruction process, albeit with a different approach on the IHCP algorithm. However, in their paper, they provide very little discussion about how the surface energy balance affects recession and environment reconstruction. They only consider melting materials, so the surface location is easily identified by the isotherm corresponding to the material melt temperature.

The current work advances the decoupling concept by incorporating the thermochemical ablation surface energy balance into the process for reconstruction of environments on materials with more complicated oxidation recession mechanisms.

III. Implementation

A code has been developed to evaluate the SEB and recession model given a set of *CHAR* solution files. While many aspects of the code focus on managing data from other codes and models, the thermochemical ablation model is the core element of this code. The model proposed by Kendall²⁵ in the ACE code (later extended by Milos and Chen²⁶) is the primary physics-based kinetics model implemented. This recession model assumes that one or a small number of heterogeneous reactions are fast relative to all of the other heterogeneous reactions but slow relative to all of the homogeneous reactions. In this scenario, the gas phase ablation products are assumed to be in chemical equilibrium, but only a portion of the char atoms (determined by the slow reactions) are able to “react” with atoms from the boundary layer and pyrolysis gas. To simplify the following discussion, it is assumed that the char is pure carbon.

Limiting the amount of char that can enter the ablation products is enabled by defining a new element of *non-reactive carbon* (denoted here by atomic symbol $C^{(nr)}$) with the same atomic weight as normal *reactive carbon* (denoted with the conventional atomic symbol C). Most carbon-containing species in the chemical equilibrium solution, such as CO and CO_2 , are defined to exist with reactive carbon and reactive carbon can only exist in gas phase. Non-reactive carbon can exist in condensed and gas phase, but only exists in gas phase as sublimation species (species containing only carbon such as $C^{(nr)}$ and $C_3^{(nr)}$). With this chemical system established, the carbon content in char is defined to be non-reactive carbon in $\tilde{\omega}_{k_c}$. The heterogeneous reaction model then converts an appropriate amount of non-reactive carbon to reactive carbon through the \tilde{B}_r' term in the elemental mass balance equation

$$\tilde{\omega}_{k_w} = \frac{\tilde{\omega}_{k_e} + B'_c \tilde{\omega}_{k_c} + B'_g \tilde{\omega}_{k_g} - B'_f \tilde{\omega}_{k_f} + \tilde{B}_r'}{1 + B'_g + B'_c - B'_f}. \quad (10)$$

where

$$B' = \frac{\dot{m}''}{\rho_e u_e C_m} = \frac{\dot{m}''}{C_m^*} \quad (11)$$

is a nondimensionalized mass flux, and $\tilde{\omega}_k$ is an elemental mass fraction. Using this mass balance, the typical diffusion limit solution approach is followed, whereby B'_c is determined such that non-reactive carbon is saturated in the gaseous ablation products. Keeping in mind that non-reactive carbon can only exist in gas phase as sublimation species, only a small amount of non-reactive carbon will be present, and B'_c will be largely driven by the carbon called for by the heterogenous reaction. This will leave the ablation products undersaturated in reactive carbon and the recession rate will be less than the diffusion limited value.

The heterogeneous reaction model for \tilde{B}_r' takes the form

$$\tilde{B}_r' = \frac{\mathcal{M}_k}{C_m^*} \sum_n \sum_i (\mu_{in}^P - \mu_{in}^R) \nu_{ki} \tilde{R}_n \quad (12)$$

where \mathcal{M}_k is the atomic weight of the limited element, μ_{in}^R is the reaction stoichiometric coefficient for reactant species i of reaction n , μ_{in}^P is the reaction stoichiometric coefficient for product species i of reaction n , ν_{ki} is the

number of atoms of element k in species i , and \vec{R}_n is the rate of reaction n . The rate of reaction for reaction n is given by

$$\vec{R}_n = K_{f_n} \left[\prod_{i,g} P_i^{\mu_{in}^R} - \frac{1}{K_{p_n}} \prod_{i,g} P_i^{\mu_{in}^P} \right] \quad (13)$$

where the products are only evaluated over gaseous species, and P_i is the partial pressure of species i . The forward rate coefficient is given by an Arrhenius model

$$K_{f_n} = B_n T^{\phi_n} \exp \left(\frac{-E_{a_n}}{\mathcal{R}T} \right) \quad (14)$$

with modeling coefficients B_n [$\text{mol}/\text{m}^2 \text{sPa}^{\mu_{\text{gas}}^R}$], ϕ_n [-], and E_{a_n} [J/mol]. The equilibrium constant is given by

$$K_{p_n} = \exp \left(- \frac{\sum_i (\mu_{in}^P \hat{g}_i - \mu_{in}^R \hat{g}_i)}{\mathcal{R}T} \right) \quad (15)$$

with \hat{g}_i is the standardized Gibbs function for species i , and the sum is over gaseous and condensed species. If the reaction is considered irreversible, the second term in Equation 13 is neglected.

Considering Equation 13, it is apparent that the reaction rate is controlled by the quantity of available reactants. The quantities of these reactants depend on the equilibrium composition and not on any limits imposed by the rate of diffusion through the boundary layer. As such, this model is a kinetic limit model and not a true transition regime model. However, it does incorporate, to some extent, the competition between reaction rate and available reactants, so it sometimes is used in place of much more complex transition regime models.

Figure 3 shows the recession predicted by the ACE kinetic model with coefficients defined to match the empirical kinetic-limit model of Scala²⁷ (both “fast” and “slow” variants). For comparison, the equivalent diffusion limit and pure kinetic limit solutions are also shown. Note that the ACE model yields the same results as the empirical models plotted with green lines when the reaction model is defined consistently. Note, however, that the ACE model departs from the empirical model as the recession rate approaches B'_c of about 10^{-2} . This is caused by the ablated carbon reducing the partial pressure of O_2 in the equilibrium ablation products mixture from the value of $P_{O_2} = 0.21$ atm used in the Scala model. Once this happens, the ACE models show similar two-plateau responses as the diffusion limit model does, albeit with transitions at increasingly higher temperatures as the heterogeneous reaction rates are decreased.

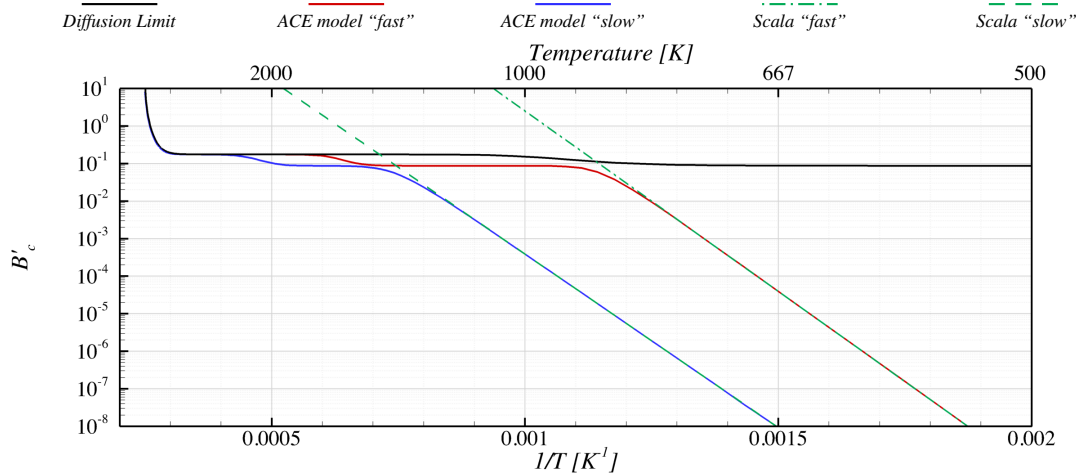


Figure 3: ACE model implementation of Scala “fast” and “slow” carbon kinetics. Nondimensionalized by $C_m^* = 0.14647 \text{ kg}/\text{m}^2 \text{s}$ and $R_b = 1 \text{ m}$.

Several attempts have been made to implement a stand-alone chemical equilibrium solver (based on^{28–30}) with the reaction model and its analytical Jacobians built into the solver. However, the chemical equilibrium problem is well known to be somewhat unstable^{28,29,31} and the algorithms implemented could not obtain solutions with enough reliability to work in this application. The multiphase Gibbs function continuation method (MPGFC) presented by

Scoggins and Magin³⁰ is purported to be completely robust in solving the chemical equilibrium problem. This algorithm fundamentally requires that the elemental composition of the mixture under consideration be constant through the solution process. Unfortunately, the elemental composition in the ACE recession model varies based on the solution species composition, which will vary through the solution.

In order to implement the ACE kinetics model with sufficient reliability to complete this work, some simplifications were required. First of all, the implementation of the MPGFC in the *Mutation++* library³² was used, with only minor modifications, to evaluate the chemical equilibrium solutions. Secondly, to get around the requirement that the elemental composition remain fixed, the solver was formulated in such a way that the *Mutation++* algorithm can be used to compute the diffusion limited recession rate for an elemental composition defined by estimated values of B'_c and \tilde{B}'_r , after which the equilibrium species composition is used to update the estimates of B'_c and \tilde{B}'_r until convergence is obtained. The process is slow but reliable, using two nested bisection-limited Newton solvers with finite-difference derivatives for B'_c and \tilde{B}'_r (refer to Oliver⁸ for more detail on the solvers); however, it limits problems to a single limiting kinetic reaction. This formulation allows a bypass of the kinetic logic to consider a simple diffusion limit recession model. Finally, this implementation presently limits problems to those with the char being composed of pure carbon.

Verification of the surface thermochemistry module was performed by code-to-code comparison and comparison to an analytic solution. For the diffusion limit case, comparison was made to the TACOT³³ B' table, which is shown in Figure 4, with excellent agreement seen^a. To verify the kinetic reactions, comparison was made to the models of Scala, and were shown previously in Figure 3, again with excellent agreement at low temperatures where it is expected to agree.

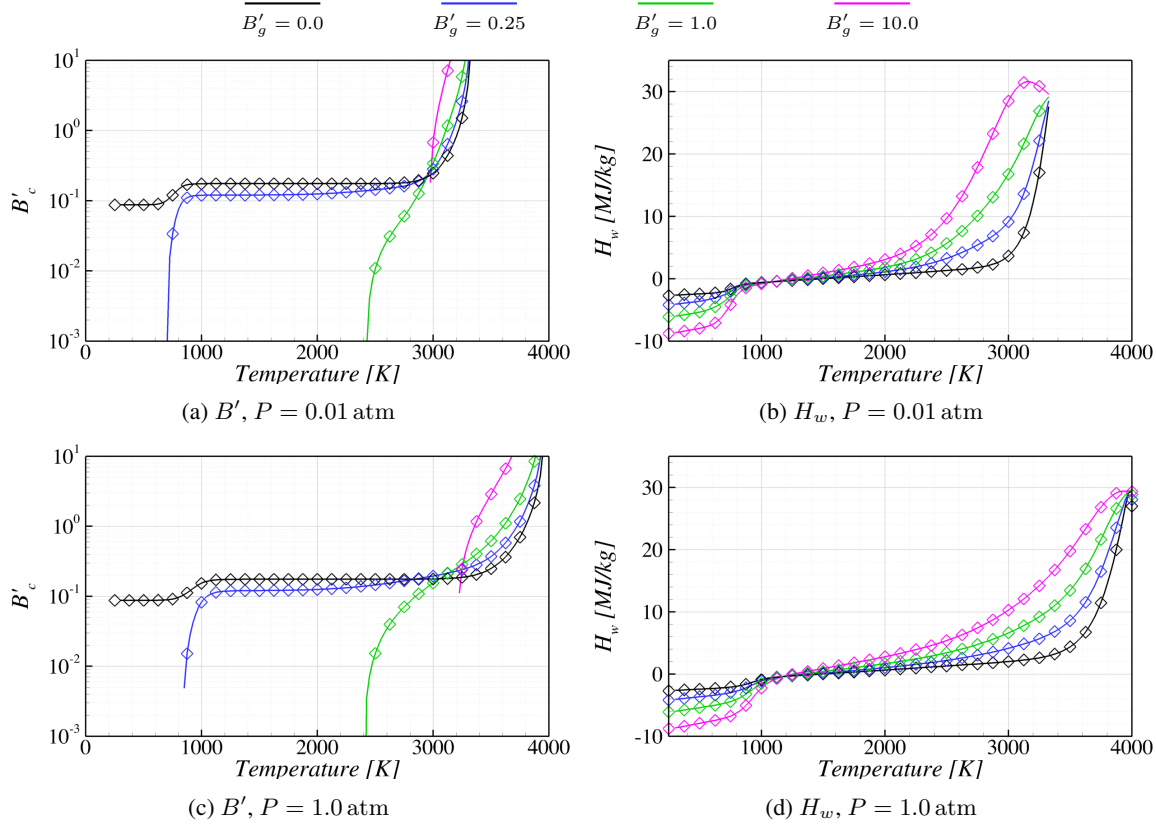


Figure 4: Verification of surface thermochemistry module by comparison to existing TACOT 3.0 B' table³³. Lines evaluated by present implementation and symbols represent published table.

^aA comment on the physics modeled in the presented table: in order to match the enthalpy at the higher B'_g values in the published TACOT table, it was necessary to allow excess carbon in the pyrolysis gas to condense. Technically this means that B'_c was negative, however the value of B'_c was simply zeroed out. This is a physical inconsistency in the published TACOT tables as the enthalpy of the condensed pyrolysis carbon is not accounted for. On the other hand, algorithms which solve for B'_c without explicitly solving for a condensed species mole number will not suffer from this, but will instead have an ablation product mixture that is over-saturated in carbon. In reality, this carbon likely would have condensed out inside the char layer through a process known as “coking”, so the real modeling failure is probably in the assumption that $\tilde{\omega}_{k_g}$ is constant.

The code is setup to compute SEB reconstructions for a range of recession model inputs in a single execution. After reading input and data structure initialization, the trajectory is reconstructed for each combination of recession model inputs specified. The presented results use an explicit algorithm in which the surface location x_s from the previous time step is used to extract the temperature, conduction flux, pressure, \dot{m}''_g , and solid density from the appropriate time solution in the temperature field data structure. The remaining material properties are evaluated at the corresponding conditions and time-dependent trajectory parameters are obtained. With that information, the explicit SEB kernel is called to find C_h^* which, along with the ablation terms determined from the specified recession model, balances the SEB. At this point, the computed C_h^* and ablation terms are written to an output buffer, the surface recession is calculated, and the loop repeats on the next time step. If the surface recession is seen to exceed a user-specified depth constraint prior to a specified time, the trajectory reconstruction is terminated. The SEB solution is obtained using an under-relaxed Newton method that seeks to sum all of the terms in Equation 4 to a tolerance of less than 10^{-8} W/m². The solver contains logic to detect and attempt to recover from an oscillating solution, and if the Newton solver fails, it will attempt to bracket the solution and converge using a bisection method. If the solver fails to find a solution, the trajectory reconstruction is terminated. An implicit algorithm has been investigated and is described further in Oliver⁸.

IV. Examples

The decoupled SEB reconstruction approach is demonstrated first with a benchmark verification then a comparison between decoupled and conventional reconstruction approaches. For each comparison, a simulated graphite test case is considered first to point out several characteristics of the method without having to consider the contribution of pyrolysis gas, and a similar TACOT³³ case is then considered to address the complications introduced by pyrolysis gas. All cases use the boundary conditions shown in Figure 5; however, the film coefficient for the TACOT case is taken to be 20% of the value shown to maintain similar recession levels on the lower density material.

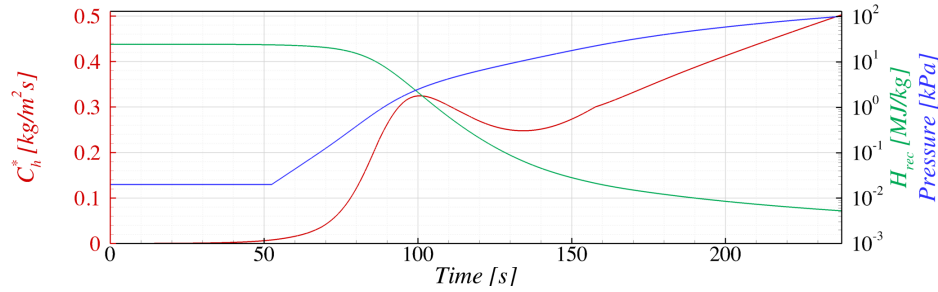


Figure 5: Boundary conditions for decoupled SEB example problems.

IV.A. Benchmark Verification

To verify the accuracy of the reconstructions obtained using the decoupled SEB method, simulated temperature measurements will be generated with high-resolution grids and recorded to 16 significant digits to reduce errors introduced by quantization. Temperature fields were reconstructed with minimal regularization, and the SEB will be evaluated to determine the reconstructed film coefficient for comparison to the film coefficient used to generate the simulated data.

An implicit assumption of this process is that *CHAR* is capable of producing simulated data consistent with the proposed physical model. Unfortunately, this is not strictly true, and in some cases the modeling fidelity of the decoupled SEB tool exceeds that of *CHAR*. *CHAR* takes the recession model in the form of a pre-computed B' table. It will interpolate in the table to obtain the surface thermochemistry solution for the current surface conditions. The decoupled SEB tool, on the other hand, directly evaluates the surface thermochemistry solution at the current surface conditions. In this regard, the decoupled SEB tool can more faithfully represent the proposed recession model than *CHAR* can. The effect of this inconsistency can be minimized by using an extremely refined B' table when generating the simulated data. Furthermore, kinetically limited recession models cannot be put into a standard table form due to the presence of the non-dimensionalizing film coefficient in Equation 12. As a result, this verification will be limited to diffusion limit recession models.

IV.A.1. Graphite

The material properties used in this case ($\rho = 1610 \text{ kg/m}^3$, $k = 8.7 \text{ W/m}\cdot\text{K}$, $C_p = 697 \text{ J/kg}\cdot\text{K}$, and $\epsilon = 0.9$, pure carbon diffusion limit recession model illustrated in Figure 6) approximate graphite, a commonly considered surface ablator. The B' table used to generate the simulated measurements was much more refined than that presented in Figure 6, containing solutions on 5 K temperature intervals at 275 different pressure levels (approximately $2 \cdot 10^5$ individual thermochemical solutions) to minimize the error introduced by interpolation in *CHAR*. The 40 mm domain is discretized with 1000 uniformly-distributed elements, and the solution is integrated through time with 0.05 s time steps. Simulated thermocouple data is extracted from a depth of 6.35 mm.

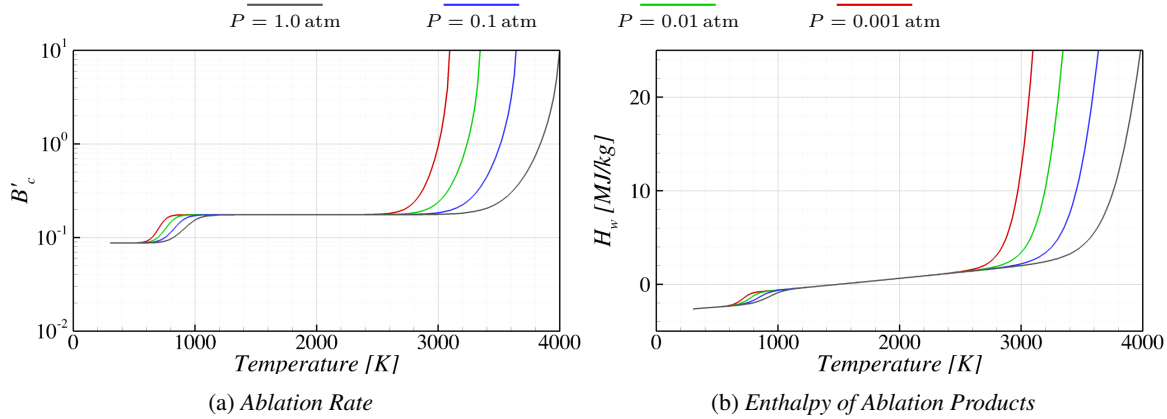


Figure 6: B' tables for pure carbon.

Reconstruction of the temperature field the non-recessing domain was performed using the SSD algorithm⁷ with a future time window of 3 s, a regularization scaling parameter $a_{-1} = -0.01$, and solution intervals of 0.05 s. Reconstruction was performed using a grid with 301 uniformly distributed elements, and the boundary of the non-recessing domain was located such that the restricted system initial surface location of $x = 0 \text{ mm}$ was one full element into the domain. Padding the reconstructed domain in this regard is done to allow the decoupled SEB tool to use a second-order central difference stencil to approximate the conduction flux on the reconstruction domain, where temperatures are defined on nodes. While the temperature field is available with a temporal resolution of 0.05 s, the SEB may be evaluated at a coarser resolution. To assess the convergence behavior of the explicit integration of the surface recession, the SEB will be reconstructed with resolutions of 0.05 s, 0.1 s, and 0.5 s.

Results of the SEB reconstruction are shown in Figure 7, with the film coefficient shown in 7(a), the surface location shown in 7(b), the enthalpy of ablation products shown in 7(c), and the aeroheating flux $\dot{q}_{aero}'' = C_h^*(H_{rec} - H_w)$ shown in 7(d). For all of these quantities, the percent error in each reconstructed value is shown and the absolute reconstructed values are shown qualitatively with the desaturated lines (reconstructed) and symbols (true values). Two qualifications on the presented results are the following: the relative error in the enthalpy of ablation products spikes as the nominal value passes through zero, and the high relative errors in surface recession at early times are due to the true value being very near zero.

It can be seen in all four plots that reducing the SEB reconstruction time step reduces the reconstruction error as expected, and the smallest time step yields very little error. Recall from Figure 6 that B'_c for this material exhibits a dual-plateau behavior, with the transition between plateau levels being pressure and temperature dependent. The peak in SEB reconstruction errors observed at 59 s appears because the temperature and pressure at this time place B'_c in this transition regime, and interpolation error was introduced by *CHAR* when simulated data was generated.

Of the four quantities presented, the film coefficient in Figure 7(a) is the primary term of interest (the others result from the film coefficient and thermal response). Given that, it is fortunate that the decoupled SEB produces reconstructions accurate to better than 0.1% for all temporal resolutions. The decoupled SEB integration yields slightly high values of the film coefficient before the film coefficient peaks at around 100 s. Recall from the definition of B' (Equation 11) that recession rate scales with film coefficient; when the film coefficient is increasing, recession rate is increasing as well. The decoupled SEB tool holds recession rate constant over the time between SEB reconstructions, with the explicit algorithm using the recession rate at the start of this interval. For a large time step when the film coefficient is increasing, the explicit algorithm will use the lowest recession rate in the integration, leading to an under-prediction of total recession. The film coefficient must then be larger to compensate for the shortfall of energy

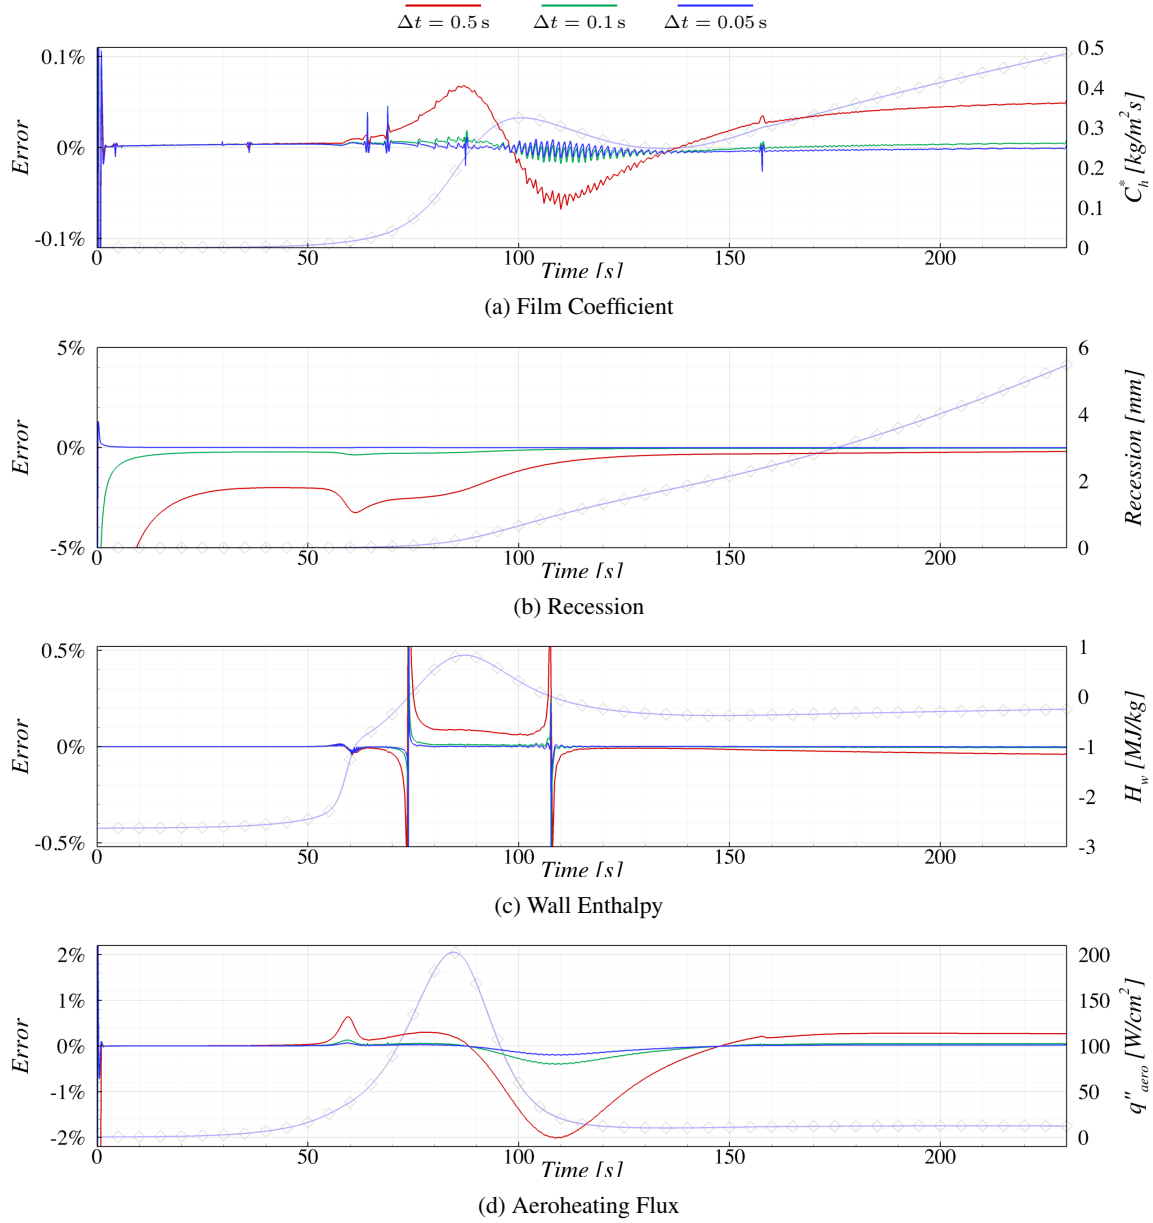


Figure 7: Percent error in reconstructions performed with decoupled SEB method on diffusion limited graphite case for a range of reconstruction time steps. Reconstructed values shown with desaturated colors and true values shown with symbols with respect to the right axis.

released by the exothermic recession reaction. The converse is true when the true film coefficient is decreasing. As the integration time step is reduced, this error is reduced as well.

The high frequency waviness that is observed in each curve is an artifact that results from a change of interpolation basis in the direct *CHAR* simulation (a compressing mesh is used, and a small discontinuity in the temperature is introduced when a grid line crosses the thermocouple extraction depth). The larger spikes are caused by small temperature discontinuities due to discontinuities in the first derivative of the true film coefficient used in the direct *CHAR* simulation. The first-order regularization used in the temperature field reconstruction will make it difficult to perfectly resolve these discontinuities, introducing the slight errors that manifest as these spikes.

IV.B. TACOT

The effect of the ablator permeability on the validity of the decoupling assumptions on the pyrolysis gas flow field were described in Section II.A. The TACOT material model has been defined with permeability values that support modest internal pressure gradients, which undermines the decoupling assumptions given that only the true surface pressure is known. However, the PICA material model for which the decoupled method is applied was built assuming negligible pyrolysis gas residence time. In this section, the TACOT model is used with both pyrolysis gas assumptions to show the validity of the method when assuming negligible residence time, as well as the potential errors introduced with more realistic physical assumptions. Furthermore, two methods of incorporating the zero residence time assumption into a 1-D *CHAR* solution is presented to determine the best strategy for reconstructions using the proposed decoupled approach.

The three permeability models considered in this example are provided in Table 1. Model A is the realistic permeability model assumed in the standard TACOT model, with the virgin permeability slightly lower than the char permeability. Model B is the model typically used when the CMA zero-residence-time assumption is used in a *CHAR* solution, with a small permeability in virgin and a large permeability in char. Model C is introduced in this work to attempt to improve decoupled SEB reconstructions and assumes that both virgin and char use the same large value for the permeability.

Table 1: Permeability models used in TACOT example problem in units of m^2

Material State	Model A	Model B	Model C
Virgin	$1.6 \cdot 10^{-11}$	$5.0 \cdot 10^{-14}$	$5.0 \cdot 10^{-9}$
Char	$2.0 \cdot 10^{-11}$	$5.0 \cdot 10^{-9}$	$5.0 \cdot 10^{-9}$

The 40 mm domain was represented by geometrically stretched elements (initial spacing 10^{-4} mm) with approximately 1500 elements. Due to the current formulation of the *CHAR* code, the pyrolysis gas mass flux is reduced to first-order at the surface, leading to the need for the highly refined mesh at the surface to appropriately resolve B'_g . To minimize error introduced by *CHAR* interpolating in the B' table, the TACOT B' table was recomputed with 139 logarithmically-distributed pressure tables, each with 244 linearly-distributed B'_g entries, each containing temperature solutions at 5 K intervals (totaling more than 25 million individual solutions). The laminar Kays^{12,34} blowing reduction model was used. The direct thermal response to generate simulated TC data was integrated using 0.05 s time steps.

The effect of the permeability model on the ablator response can be seen in Figure 8. Each of the three permeability models considered are shown, along with their corresponding non-dimensionalized blowing rates. Models A and C, with relatively high virgin permeability show an earlier rise in blowing rate as compared to model B, which has a low virgin permeability. The blowing rate for model B rejoins the predictions from the other models once the surface material has decomposed sufficiently to increase the permeability to the point that gas may freely flow from the surface. Since surface recession models are typically strongly dependent on the pyrolysis gas blowing rate, this discrepancy can be expected to affect the reconstruction problem and motivated the model C formulation.

To assess the performance of the three different permeability models in the decoupled SEB reconstruction approach, the simulated TC data at 6.35 mm produced by permeability model A are used as the reconstruction target. Whole domain reconstructions⁷ of the temperature field are performed with solution intervals of 0.5 s and a regularization scaling parameter of $\alpha_{-1} = -0.001$. For the temperature field reconstruction, the domain is discretized using 900 uniformly-spaced elements, with the non-recessing boundary located such that first few elements of the domain are located outside of the initial surface to permit central-difference reconstruction of the conduction flux and to avoid numerical artifacts of the mass flux being reduced to first order at the boundary.

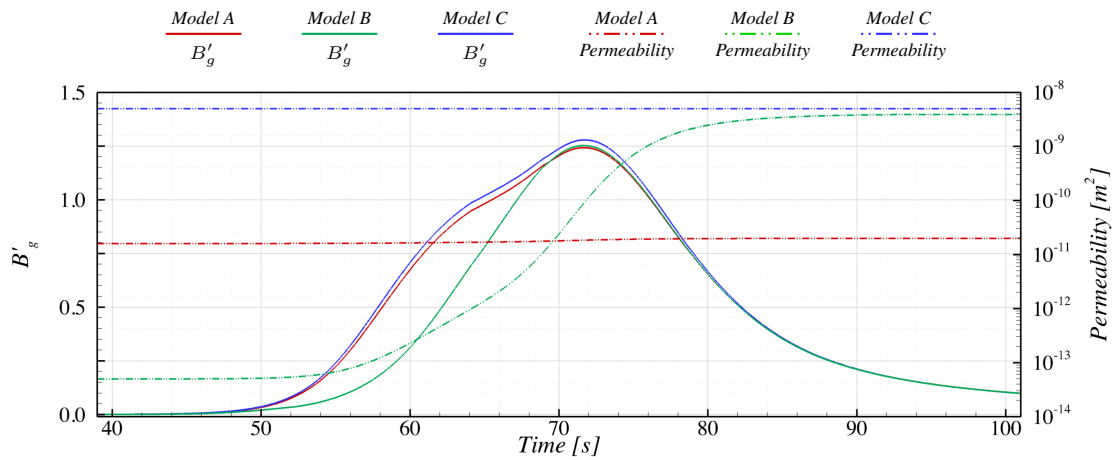


Figure 8: Pyrolysis gas blowing rate and permeability of the direct problem.

Figure 9 shows the reconstructed values (desaturated colored lines) of film coefficient (9(a)), surface recession (9(b)), aeroheating flux (9(c)), pyrolysis gas blowing rate (9(d)), and surface temperature (9(e)) with the true values plotted with symbols. The percent error of each reconstruction is shown using saturated colored lines. The red, green, and blue lines denote results utilizing thermochemical solutions computed at the local conditions, whereas the magenta dashed line utilizes thermochemical solutions in the highly-refined B' table used in the direct problem. The thermochemical solution algorithm at early times where temperatures are very low is not robust enough to yield a solution to the SEB. To work around this issue, the reconstructions are started at 30 s assuming that no recession has occurred to that point. The reconstruction utilizing the B' table did not suffer from this limitation, so it is started at 0 s. This is the source of the large errors noted in surface recession in Figure 9(b), as a very small amount of recession occurs early in the direct problem. Once significant recession begins around 80 s, the reconstructions quickly come back in line with the true values. The difference in recession does not significantly affect the reconstructed film coefficient values, as can be seen by the close agreement between magenta and red lines in Figure 9(a).

For each permeability model, over- or under-prediction of the film coefficient or aeroheating flux tends to correlate with similar over- or under-prediction of the reconstructed surface temperature. Since the general trend is to over-predict prior to peak heating and under-predict after peak heating, the errors are likely the result of the reconstructed temperature field leading the true temperature field due to the long recession integration interval (0.5 s).

The reconstructed pyrolysis gas mass flux in Figure 9(d) does not seem to show the same trends noted before, as Model C performs generally better over the time considered and Model A shows some notable (4%) error prior to 100 s. The improved prediction of blowing rate does not lead to a more accurate reconstruction of the film coefficient. Given the strong influence of blowing on the recession rate, this is a surprising but misleading result. Considering the plots of pyrolysis mass flux for each permeability model in Figure 10, the effects of reducing the mass flux to first-order at the boundary are apparent by the discontinuous behavior of the reconstructed flux contour lines (black lines) in the first two elements, and the effect is seen to decrease with increasing permeability. Though not clearly visible, these discontinuities are also present in the higher-resolution direct solutions (colored lines). While the reconstructed contour lines closely match the trend of the true flux contours extrapolated through the true system boundary elements, the actual blowing rate considered in the direct problem recession model evaluation is the rate affected by these discontinuities. This means that the blowing rate used as a reference in Figure 9(d) contains a numerical artifact not present in the decoupled reconstructions that leads to the misleading conclusion.

Recall from the discussion earlier in this section that the direct problem required approximately 1500 geometrically-distributed elements, whereas the reconstruction was performed on 900 uniformly-distributed elements. To assess the adequacy of this assumption, a grid refinement study is performed on reconstructions of permeability model A TC data. The reconstruction obtained using four increasingly refined uniformly-distributed element grids is presented in Figure 11. The introduction of the pyrolysis gas mass flux field increases the necessary grid resolution for this example to at least 600 uniformly-spaced elements, and the present resolution is sufficiently grid converged.

The performance of the three different permeability models in this reconstruction problem varies and there is no clear ‘best’ model. It is clear, however, that using the proposed Model C to improve the pyrolysis gas decoupling does not necessarily lead to better overall reconstructions. The SEB is highly non-linear, so the errors introduced by the temperature field reconstruction and errors introduced by errors in the pyrolysis gas blowing rate combine to yield

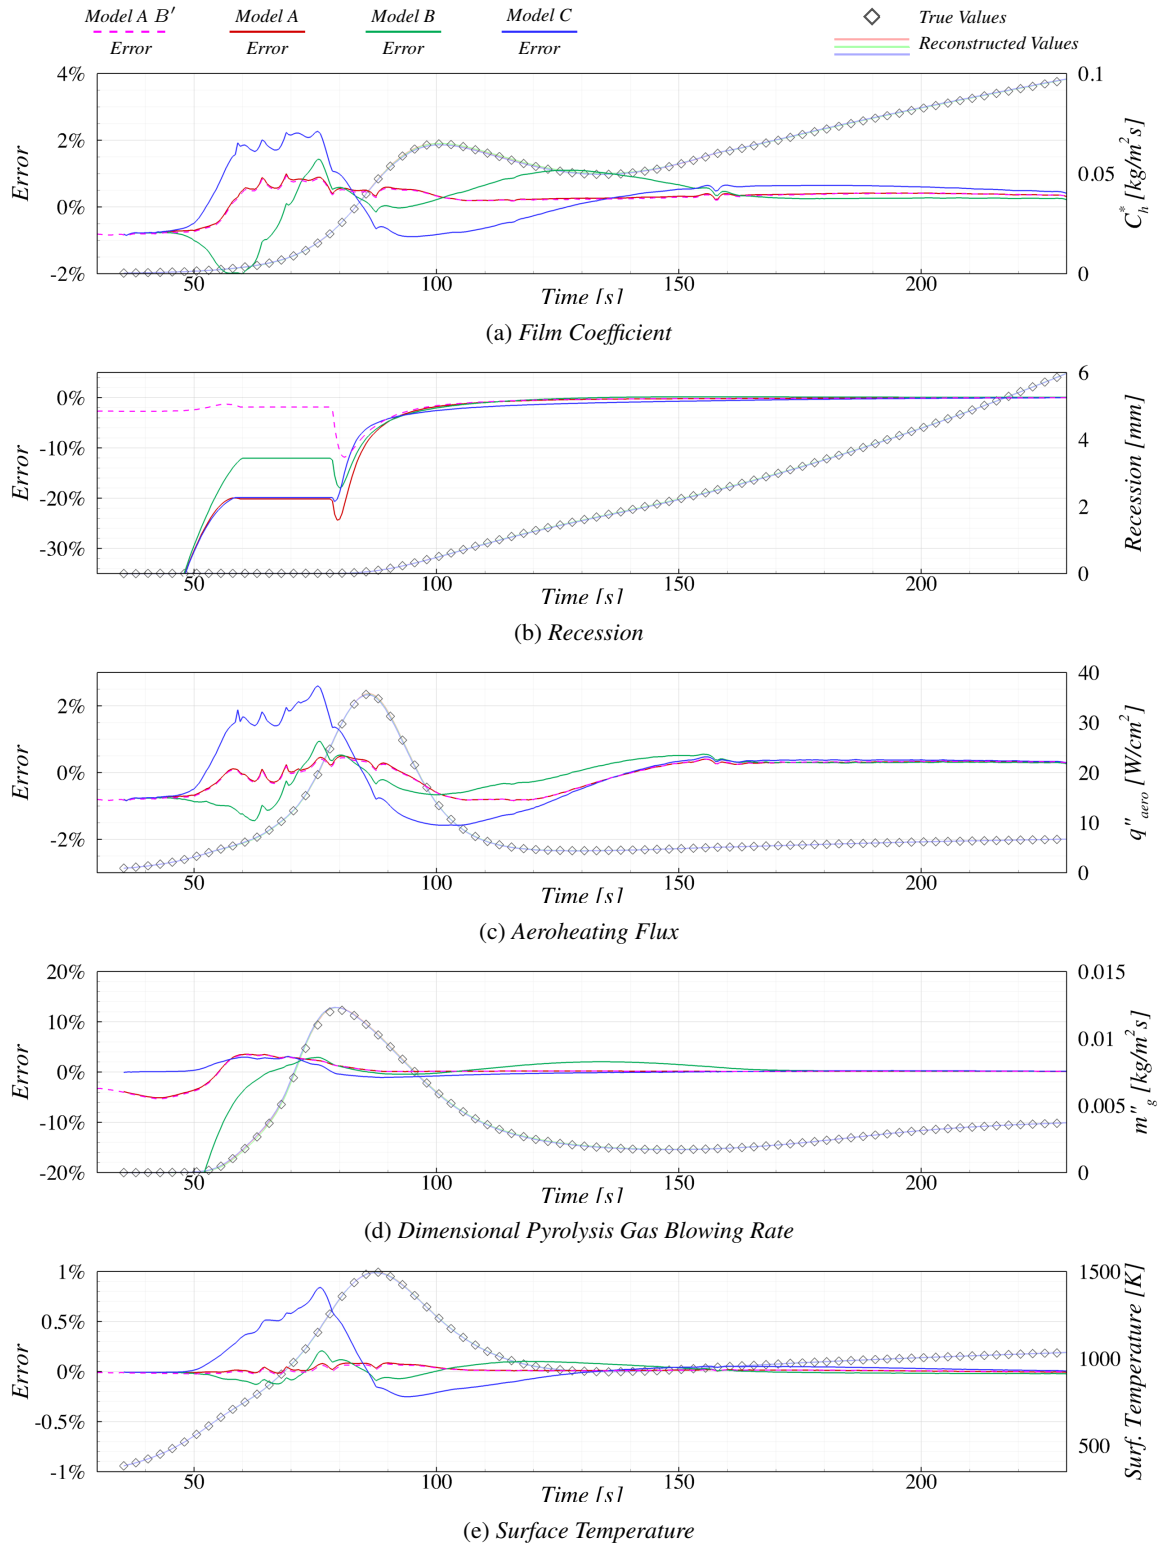
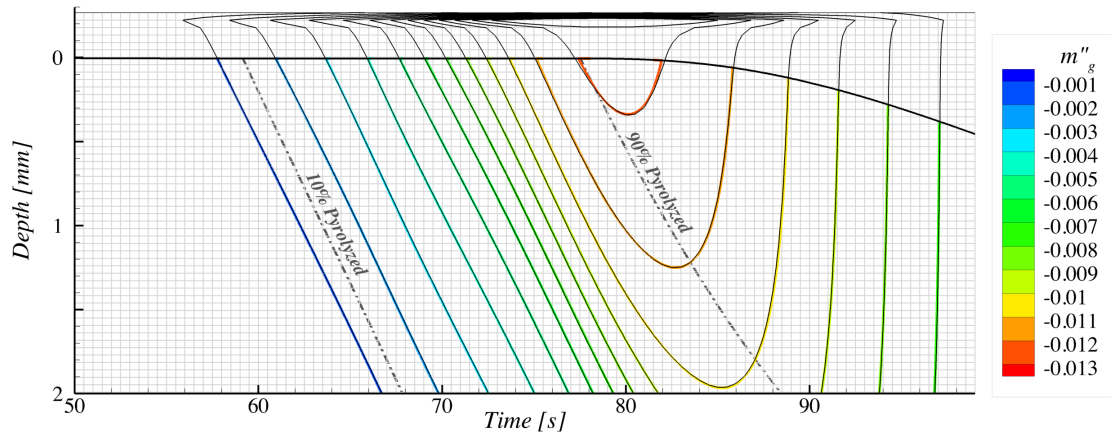
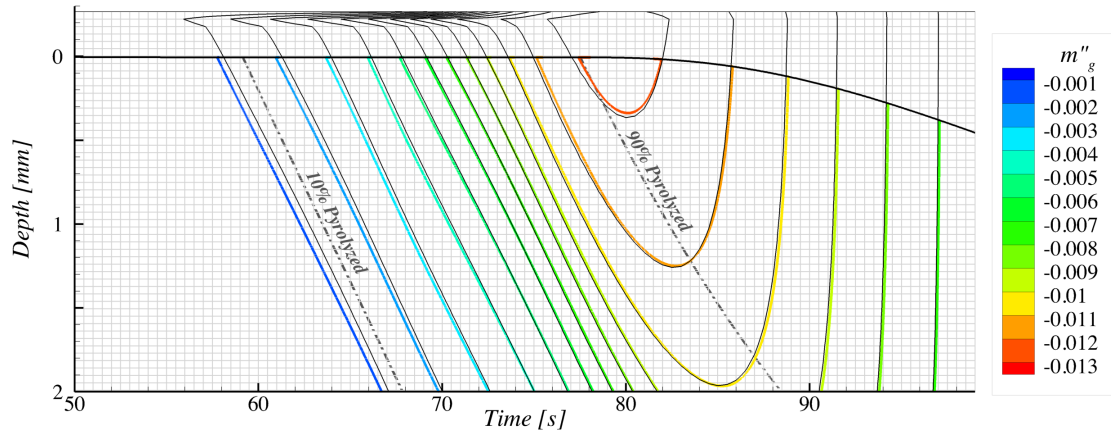


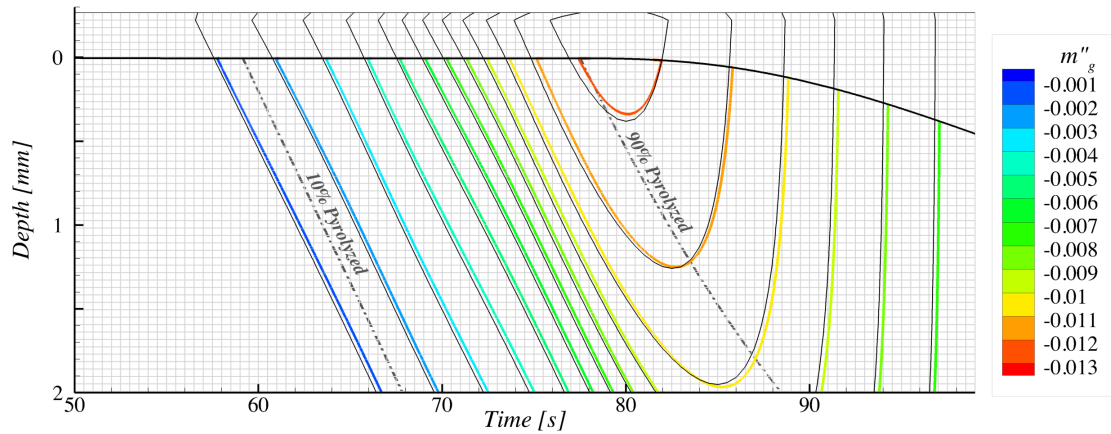
Figure 9: Percent error in reconstructions performed with the explicit decoupled SEB method on diffusion limited TACOT case for each permeability model. The Model A permeability model is used to generate the simulated data and provide the basis for the error calculation. The curve denoted with B' interpolates in a highly-refined B' table for the thermochemical solutions, the remainder of the reconstructions evaluate the solution at each point.



(a) Permeability Model A



(b) Permeability Model B



(c) Permeability Model C

Figure 10: Reconstructed pyrolysis gas mass flux field (black lines) for each permeability model overlaid on the true flux field (colored lines) for permeability model A. Grid indicates reconstructed times and mesh density, thick black line indicates true recessed surface location, and dash-dot line denotes extend of pyrolysis zone in true domain.

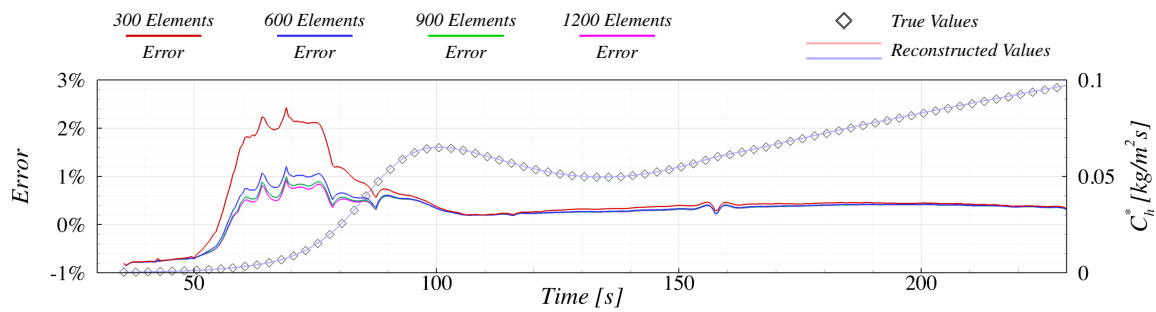


Figure 11: Comparison of permeability model A reconstructions on four levels of grid refinement. The Model A permeability model is used to generate the simulated data and provide the basis for error calculation.

the net error in the film coefficient. However, these errors in the final result are small; none of the reconstructions considered exceeds 3% (1% if Model C is excluded).

IV.C. Comparison to Film Coefficient Reconstruction

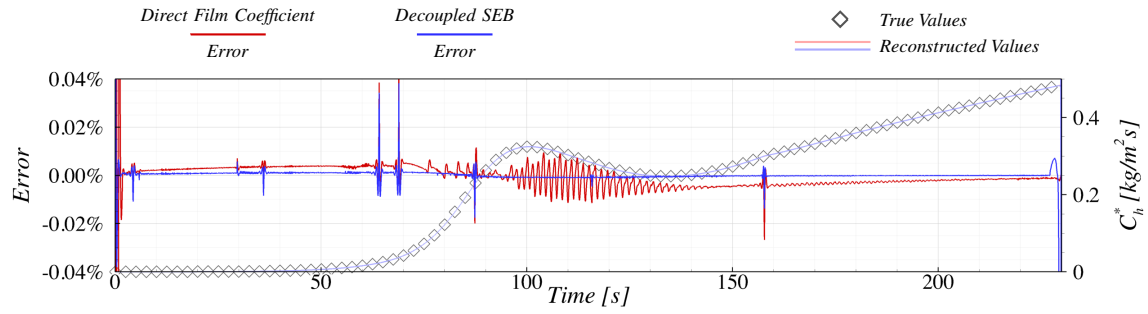
The proposed decoupled SEB reconstruction approach is an alternative to performing an inverse reconstruction directly for the film coefficient, with the recession model evaluation involved in the sensitivity coefficient evaluation. For recession models provided in B' table form, *CHAR/INHEAT* is capable of performing a reconstruction in this manner; however, the decoupled reconstruction approach provides a number of advantages for some problems. These advantages are discussed in this section as these two approaches are compared on the graphite and TACOT examples from the previous section.

In order to isolate the differences between the two reconstruction approaches, the direct reconstruction and the decoupled temperature field reconstruction are performed with the same IHCP algorithm and regularization parameters. The film coefficients for both the graphite and TACOT cases are reconstructed on the same intervals used in the decoupled SEB reconstructions (graphite: 0.05 s, TACOT: 0.5 s). Both cases use a regularization scaling parameter of $\alpha_{-1} = -0.01$ (the decoupled SEB result presented in this section was recomputed to use the same SSD algorithm as the direct film coefficient reconstruction). TACOT reconstructions use permeability Model A.

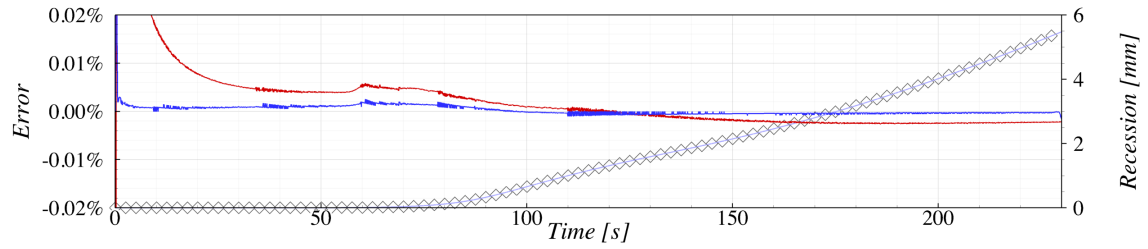
Although the IHCP parameters are consistent, the TACOT temperature field reconstruction for the decoupled SEB method is performed on the 900 element uniform mesh described before, whereas the direct film coefficient reconstruction must use the 1500 element stretched mesh since the pyrolysis gas mass flux at the surface must be resolved. This contributed to an increase in the computational cost of the direct film coefficient reconstruction. Running with 12 threads on 12 cores (Thinkmate VSX R5 760V3 workstation with 2 Intel XeonTM E5-2697 v3 2.60 GHz CPUs for a total of 28 cores with 128 GB of DDR4 2133 MHz ECC RAM), the decoupled SEB temperature field reconstruction required 3.0 hr of wall time and the direct film coefficient reconstruction required 5.8 hr. The difference in mesh size certainly contributes to this difference. However, the direct problem solutions for the decoupled SEB temperature field reconstructions are considerably cheaper due to the linearity of the boundary conditions, the boundary condition not needing to iterate on the blowing correction, and mesh motion not needing to be considered. The final step in the decoupled SEB process, integrating the recession model, required less than 4 s to evaluate on a single core. If multiple recession models are to be considered in a sensitivity analysis using the direct film coefficient reconstruction approach, the 5.8 hr reconstruction would have to be repeated for each evaluation. For the decoupled SEB method, however, only the 4 s SEB evaluation would need to be repeated, as the reconstructed temperature field does not depend on the recession model. Even neglecting the modest speedup in the temperature field reconstruction, the considerably reduced cost of the SEB evaluation enables much more detailed assessment of recession model sensitivities.

To keep the recession models on similar terms for comparison, the decoupled SEB results are evaluated using the same high-density B' tables used by the direct film coefficient reconstructions. Figure 12 shows the comparison of the reconstructed surface conditions for the graphite example. While generally very similar with little reconstruction error, the direct film coefficient reconstruction results are less accurate than the decoupled SEB results, likely due to the fact that the regularization is directly applied to the film coefficient, which impacts the other parameters through the recession model.

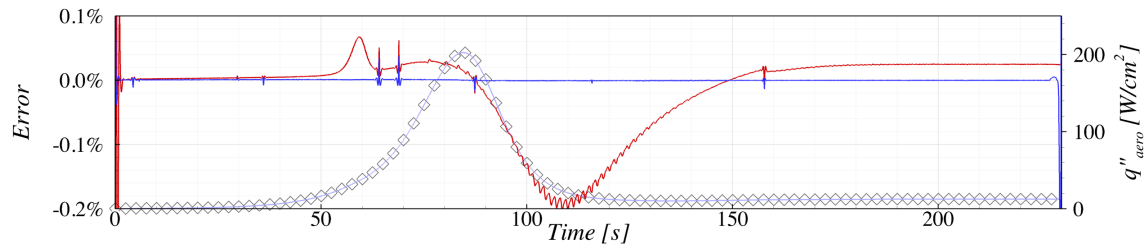
Figure 13 shows the comparison of the reconstructed surface conditions for the TACOT example. Both algorithms have large errors at the very beginning of the problem when the true film coefficient is small, but have come to reasonable values before the heating rate becomes significant. The direct film coefficient reconstruction performs



(a) Film Coefficient



(b) Recession



(c) Aeroheating Flux

Figure 12: Comparison of reconstruction method results on the diffusion limited graphite case.

markedly better in computing the pyrolysis gas blowing rate and surface recession. Presented this way, the spike in recession error around 80 s for the decoupled SEB method appears likely due to delayed onset of recession caused by the relative over-prediction of surface blowing (recall the discussion from the previous section regarding the surface mass flux in a direct *CHAR* solution). The improved accuracy in blowing and recession for the direct film coefficient method does not translate to more accurate film coefficients or aeroheating fluxes, with the latter two quantities being more relevant for flight vehicle design. The decoupled SEB method tends to provide a more consistently accurate estimate of the aeroheating flux. In terms of the film coefficient estimate, neither method is obviously better than the other, and errors are generally less than $\pm 2\%$. This indicates that even with the potential errors by decoupling the pyrolysis gas mass flux field with realistic permeability values, the decoupled SEB method provides reconstructions of generally equivalent accuracy at a greatly reduced cost.

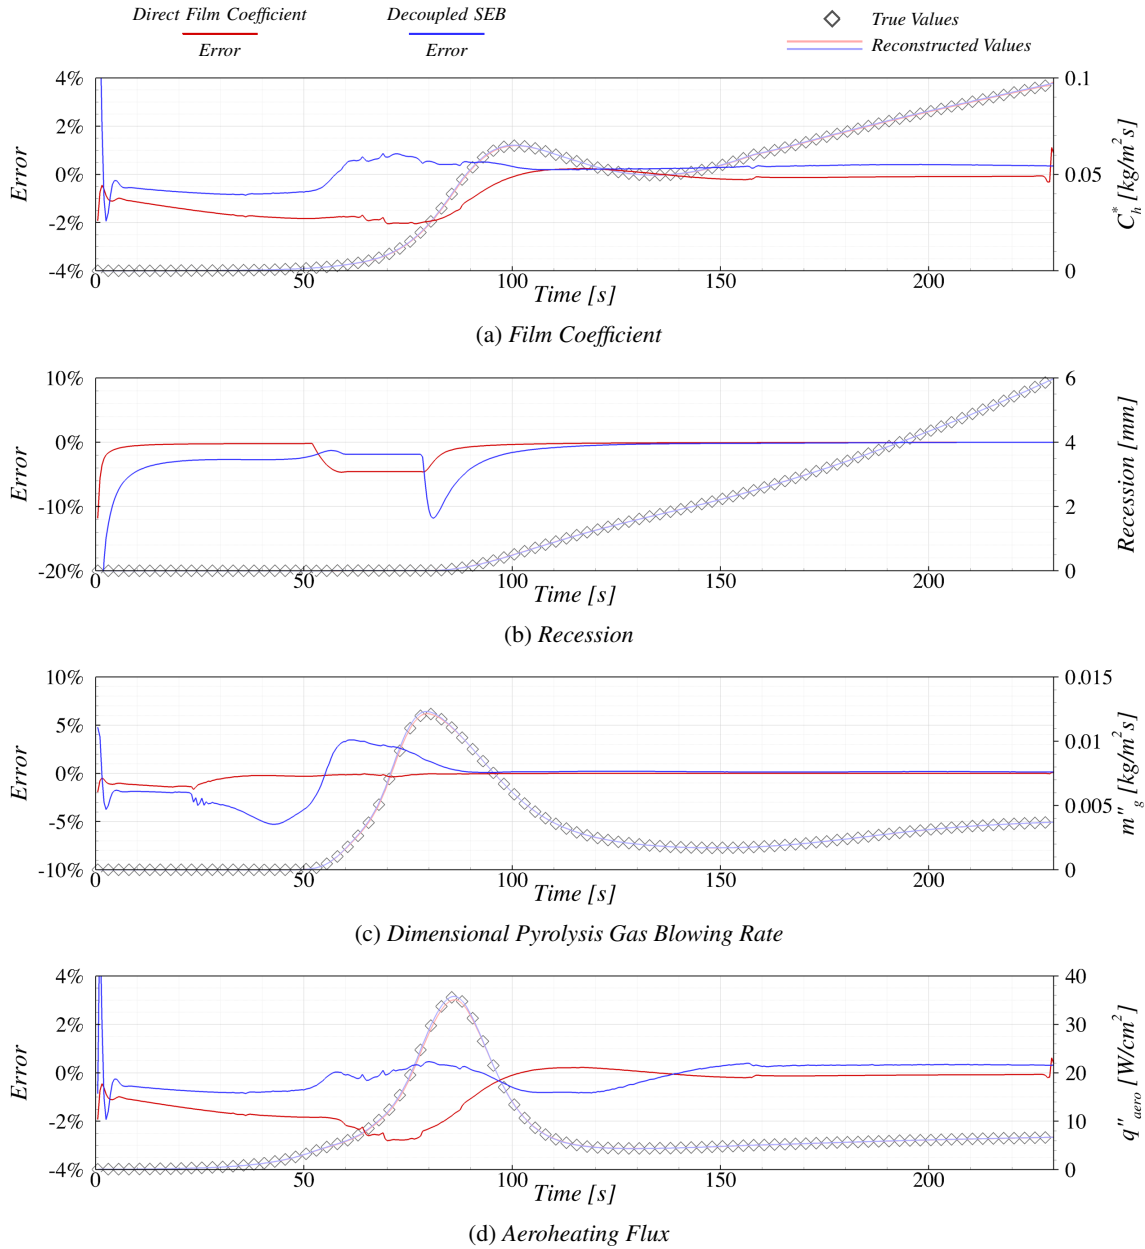


Figure 13: Comparison of reconstruction method results on the diffusion limited TACOT case.

V. MEDLI Reconstruction

The PICA heatshield of the MSL capsule was fitted with seven instrumented plugs to gather temperature data during its August 2012 entry into the Martian atmosphere^{2,3,5}. Each contained four TCs embedded within the PICA plug, the shallowest nominally 0.1 in from the un-ablated surface. An unexpected result of the flight was that none of the TCs failed during the entry, suggesting that the surface recession was less than 0.1 in. This indicates that the surface recession model used in design and post-flight analysis is notably over-conservative since it predicted much more recession. Curiously, it was also shown by Bose et al.³ that predicted heating environments and the PICA thermal response model together under-estimate the temperatures near the stagnation point, with the predicted integrated heat load 33% below the flight heat load determined by Mazhari et al^{20,23} using inverse methods. Cruden et al.³⁵ theorized that shock-layer radiation, which was not accounted for in the predicted environments, could be contributing to the discrepancy. Consequently, they performed some testing and analysis activities to generate shock-layer radiation estimates for the flight and show that the discrepancy can be reduced to nearly half by including radiation estimates in the direct thermal analysis. They propose several different radiation profiles combined with different recession assumptions to show the temperature and load predictions improve, but it is unclear which profile is ‘right’ and what other models used in the direct analysis could be contributing to the remaining discrepancy^b.

This section applies the techniques presented in this paper to two of the MEDLI thermocouple plugs. The SEB reconstruction results for a number of recession and environment assumptions are discussed. In some instances, models are shown to provide infeasible results thereby allowing some conclusions to be made regarding the applicability of certain assumptions. That said, it should be cautioned that this approach is not guaranteed to provide clear statements on the adequacy of specific models. The environment and surface thermochemistry form a tightly-coupled system, and the interdependence of each component means that no single aspect can be clearly addressed by one measure of the overall response. This is an unfortunate example of the possible non-uniqueness in the inverse problem. The proposed approach does, however, allow viewing the system response in a different light, providing another measure of the performance of each part of the complicated mechanism.

V.A. Flight Information

Temperature data from the nearest-surface thermocouples in MISP 1 (near the stagnation point) and MISP 7 (observed the highest temperatures and clear laminar/turbulent transition) are used for reconstruction. Pressure ports MEADS 2 and MEADS 5 are assumed to be close enough to MISP 1 and MISP 7, respectively, that the measured surface pressures are used in the reconstruction without spatial interpolation. The recovery enthalpy was taken to be the total enthalpy calculated using the altitude, velocity, and atmospheric properties reconstructed from MEADS and IMU (inertial measurement unit) data.

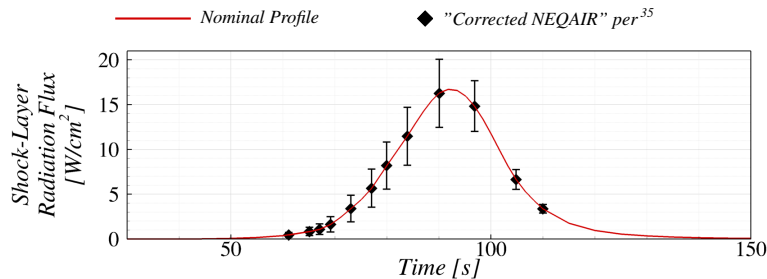


Figure 14: Shock layer radiation used for reconstruction of MISP data.

The shock-layer radiation profile and surface emissivity are not well understood and are treated stochastically in a sensitivity analysis. Figure 14 shows the shock-tube data corrected prediction of shock-layer radiation heating on the MSL stagnation point from Cruden et al.³⁵ along with uncertainty estimates. A spline-fit curve of the nominal test data is used as the baseline, and each sample of the sensitivity analysis uses the baseline multiplied by a perturbation

^bThis highlights a key limitation of so-called *direct comparisons* to flight data. Especially on ablators, many models must be used in concert before a comparison may be made to available data. An error in a model applied early in the process could drive an otherwise accurate model applied later in the process to provide the wrong result. As a thought experiment, imagine the case of under-predicted environments that drive surface temperatures to a low enough value that an accurate temperature-dependent recession model does not predict enough recession. It would be easy to suggest that the recession model is inaccurate, especially if the environment prediction were difficult and costly (thereby implying a higher level of accuracy than may be truthfully warranted). If the recession model could have been applied with surface temperatures more representative of reality, it would be allowed to demonstrate its accuracy and the fault in the environments could more accurately be identified.

drawn from a uniform distribution on the interval $[0.78, 1.22]$. In arcjet testing, the RTV adhesive around the plugs was shown to deposit silica on the face of the plugs³⁶ at low heating conditions, so the surface emissivity was perturbed by uniformly-distributed factors of $[0.75, 1.0]$ to address this uncertainty.

The temperature data in reference² is provided with a resolution of 0.1°C . Aside from this quantization error, there is very little evidence of high-frequency noise in the TC data. As a result, the TC data was not filtered prior to the reconstruction.

V.B. Temperature Field Reconstruction

The boundary conditions were reconstructed on the 1-D as-designed stack up using the SSD algorithm⁷ at frequencies for which data is available (8 Hz for MISP 1, and 2 Hz for MISP 7), with the reconstruction future time window taken to be 5 s long. Time was integrated with 0.0625 s time steps from a uniform initial condition that matched the temperature measurement at entry interface. A grid refinement study showed that a grid consisting of 800 elements in the PICA layer was sufficient⁸.

Locally-scaled first-order regularization factors were used, with four scaling parameters from 1.0 to 0.001 considered. Figure 15(a) shows the reconstructed temperatures at peak heating and Figure 15(c) shows the difference between the reconstructed and measured temperatures for MISP 1. Results using regularization scaling parameter $a_{.1} = 1.0$ visibly under-shoot the peak temperature in 15(a) and also demonstrate visible bias relative to the data and other reconstructions in 15(c). By contrast, results for the next largest value of the scaling parameter do not show these traits, therefore this value ($a_{.1} = 0.1$) is used for the presented results. Since the data rate is lower for MISP 7, a regularization parameter of $a_{.1} = 0.01$ performs better on that plug (Figures 15(b) and 15(d) shows the equivalent data for MISP 7).

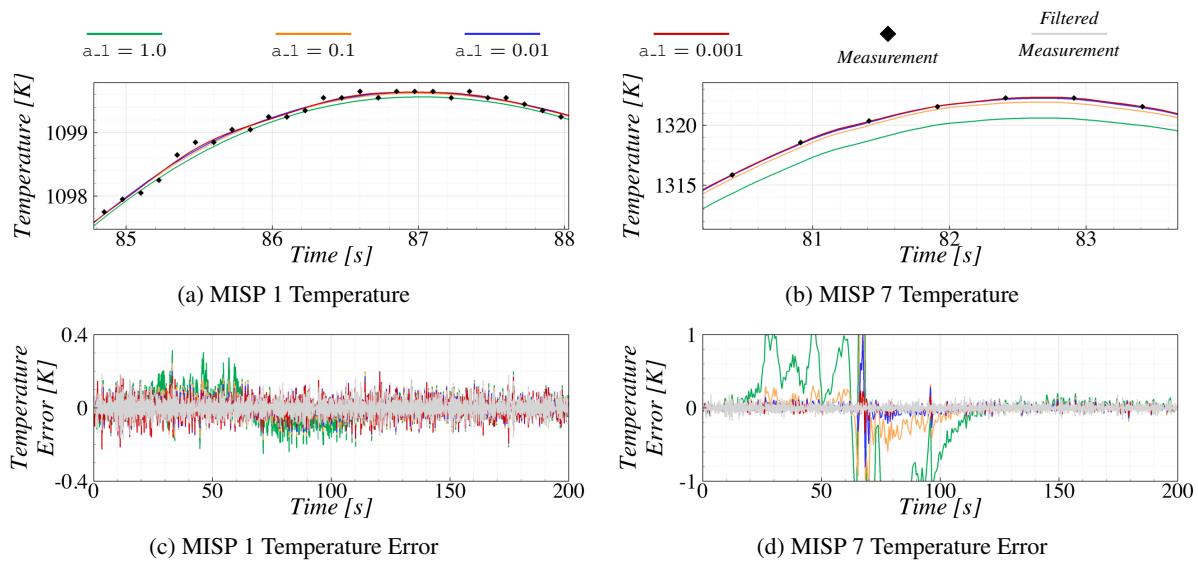


Figure 15: Reconstructed temperature and residual temperature error at shallow TC using four different values of the first-order regularization scaling factor.

Reconstructed surface temperatures (neglecting recession) are shown in Figure 16, along with temperatures at the first two TC locations. As expected given the measurements at the TCs, the surface temperature for MISP 7 is quite a bit higher than that for MISP 1. The sharp temperature increase at approximately 65 s (that can be attributed to higher heating due to laminar to turbulent boundary layer transition) on MISP 7 is much more pronounced at the surface compared to the first TC depth, which is used as the reconstruction target. Furthermore, the good comparison of the modeled and measured temperatures at the second TC (blue line) show the relatively high-quality of the conduction and decomposition components of the PICA thermal response model, since measurements from the second TC did not influence the reconstructed boundary conditions.

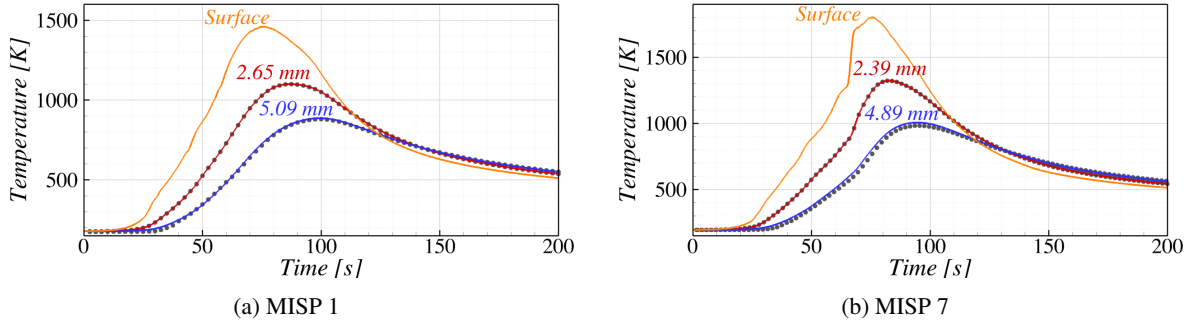


Figure 16: Reconstructed surface temperatures and TC temperatures relative to measured data. A subset of the measured data is denoted by filled circles. Recall that the shallow TC (red line) is used as the reconstruction target.

V.C. SEB Reconstruction

Decoupled SEB reconstructions of MISIP 1 highlight a useful capability of decoupled SEB sensitivity analysis; namely, that “other data” and the existence of a solution to the SEB can be used to gain insight into the appropriateness of the modeling inputs. In this instance, our “other data” is the knowledge that recession did not exceed 0.1 in. Any inputs that result in an SEB reconstruction showing more recession than this are marked as *infeasible*. Furthermore, if the SEB cannot be balanced while respecting modeling assumptions (primarily the assumption of a positive film coefficient), then the solution is terminated and the modeling inputs are marked *infeasible*. Inputs that remain within observed constraints and are consistent with the modeling framework are marked as *feasible*.

Two recession models will be considered in the sensitivity study: a scaled diffusion limit model, and a one-reaction CO_2 model. The scaled diffusion limit model assumes that at each point in the reconstruction, the actual B'_c value is a constant factor (the model input) of the diffusion limit B'_c evaluated at those conditions (the scaled B'_c is then used to evaluate the wall enthalpy assuming gas-phase chemical equilibrium). There is no physical basis for this scaled diffusion limit recession model, so a simple CO_2 model has been formed for this study^c based on the data of Gulbransen et al.³⁷. The reaction is assumed to have the form $CO_2 + C(s) \rightarrow 2CO$, and the model parameters B_n and $E_{a,n}$ are stochastically defined contingent on approximately fitting the experimental data (Figure 17).

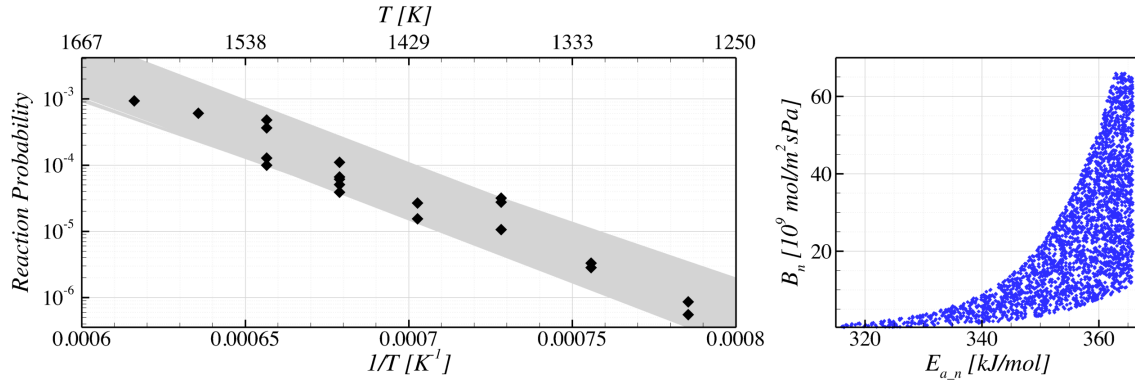


Figure 17: Range of model inputs considered in sensitivity study and corresponding reaction probabilities relative to test data of Gulbransen et al.³⁷.

Figure 18 shows feasibility scatter plots and histograms for the stochastic variables considered in the SEB sensitivity study of MISIP 1. The red symbols denote infeasible input combinations and the blue symbols denote feasible input combinations. Based on Figure 18(b), it appears that there is a strong suggestion that the B'_c scale factor must be less than 0.2, otherwise the surface reconstruction yielded too much recession. Furthermore, it is seen in Figure 18(a), which includes results from both recession models, that the shock layer radiation scale factor cannot exceed 1.06 (with a slight dependence on the emissivity scale factor) without driving the solutions into infeasible territory. In this

^cSee reference⁸ for derivation details, though note that the CO_2 model SEB reconstruction results presented here corrects an error identified in that work.

latter case, the radiation provides enough heat into the surface that the boundary layer must produce a negative flux to balance the SEB and remain consistent with the thermocouple observations. When the radiation factor is too large, this cooling is necessary at a point when the recovery enthalpy is still greater than the wall enthalpy requiring the film coefficient to be negative; a condition not permitted in the present framework. Note that in this study, approximately 15 thousand samples were evaluated using the scaled diffusion limit model, but only 3.7 thousand samples were taken using the CO_2 model due to cost. The scaled diffusion limit solutions take less than 8 s on a single core, whereas the kinetically limited solutions (denoted in the Figure 18(a) histogram with darker colors) can take up to 8 hours each.

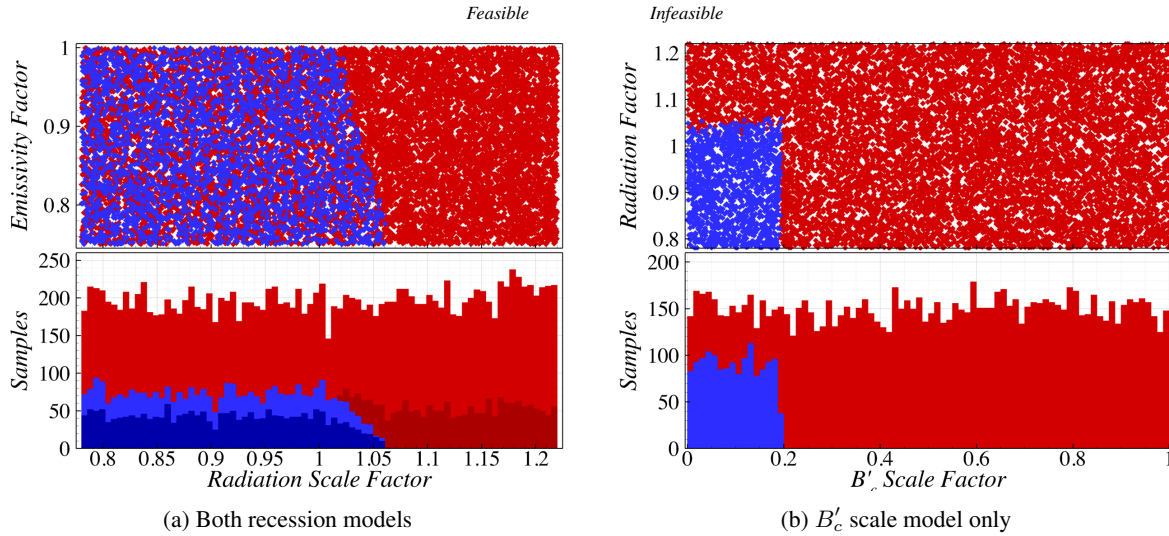


Figure 18: Feasible (blue) and infeasible (red) samples for MISP 1 sensitivity study .

The observation that there is an upper bound on the shock layer radiation is a significant one, as this paves the way for a reduction in the shock layer radiation prediction uncertainties. Before this reduction can be confidently applied, the sensitivity study must be grown to include parameters that affect the temperature field reconstruction (for instance, the char thermal conductivity or heat of pyrolysis). This study will be notably more expensive to perform since the SST temperature field reconstruction for MISP 1 required approximately 1200 CPU-hr, whereas each scaled diffusion limit SEB reconstruction required^d less than 8 CPU-hr.

The higher heating seen on MISP 7 did not drive SEB solutions into infeasible territory when the radiation scale factor was too large; however, it is not expected that radiation will be higher at this plug location. Consequently, only samples that fell in the feasible space indicated in the scatter plot of Figure 18(a) are considered in the MISP 7 reconstruction.

Reconstructed environments on MISP 1 and 7 are shown in Figures 19 and 20 respectively, with the Kays laminar blowing reduction model used to obtain the unblown film coefficient. The lines represent the mean the feasible solutions at each point in time, and the error bars show ± 1 standard deviation. The orange lines denote the CO_2 kinetic recession model assumption, the blue line denote the assumption of no recession, and the green lines denote diffusion limited recession. Since diffusion limit is a poor assumption in this case these solutions become ill-behaved at later times; hence these curves have been truncated at 85 s. It is interesting to note the behavior of the kinetically limited solutions relative to the bounding no-recession and diffusion limited results. Prior to peak heating, the kinetically limited solutions very closely follow the no-recession bound. Once the surface temperature has sufficiently increased to drive the modeled heterogeneous reaction, recession picks up and in the case of MISP 7 closely approaches the diffusion limited rate. Once temperatures cool after peak heating, recession begins to drop off. Differences in surface location and sensitivity to small differences between the wall enthalpy and the recovery enthalpy lead to larger differences between the kinetically limited and no-recession reconstructed film coefficients at late times.

Considering the additional heat released by the surface recession reactions, it makes sense that the higher-recession models yield a lower aeroheating flux. Recall that for a reconstruction, the heat absorbed by the heatshield is fixed, so the boundary layer provides less heating if more heating is provided by recession. In a CO_2 atmosphere, adding carbon to the boundary layer gas results in a higher wall enthalpy, so the film coefficient may either increase or decrease, depending on which effect is stronger.

^dNote that temperature field reconstruction was performed with $\Delta t = 0.125$ s whereas the SEB reconstruction was performed on a $\Delta t = 0.25$ s.

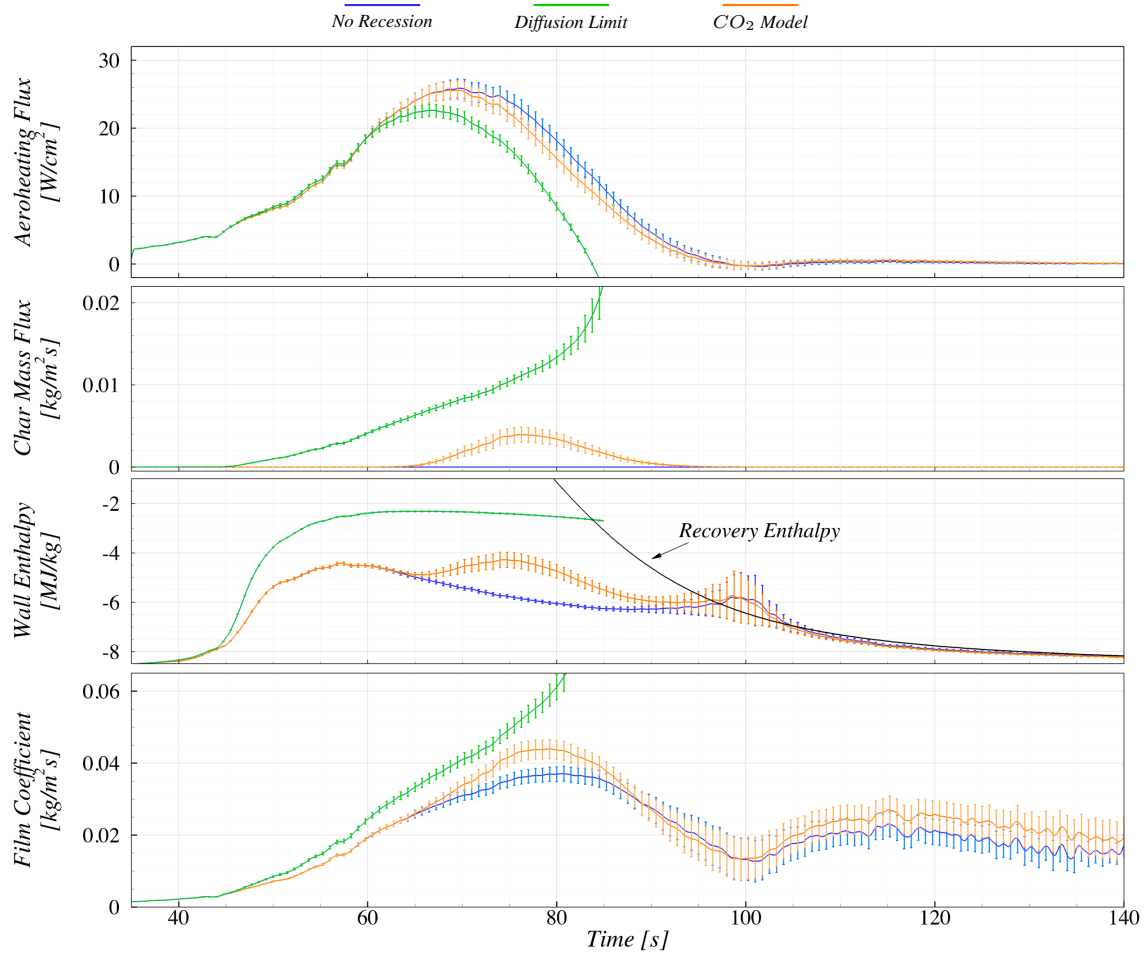


Figure 19: Reconstructed environments on MISP 1.

VI. Conclusions and Future Work

In this paper, the theory for the decoupled SEB reconstruction approach is presented and discussed. The role of the ablator permeability on the internal pressure field is discussed and it is described how this could undermine the validity of a decoupled reconstruction. A description of the code and recession models implemented to execute the proposed algorithm is provided, and several examples are presented. The first shows that the decoupled SEB algorithm is capable of producing very accurate reconstructed film coefficients (better than 0.1%) on materials without internal decomposition and pyrolysis gas. Incorporating pyrolysis gas decreases the accuracy of the reconstructions, although the film coefficient errors remained less than 3%. Several different models for representing the low-density ablator permeability are assessed with the apparent result that there is no justifiable reason to use an artificial permeability in the decoupled reconstruction process if a physically-accurate value is available. The decoupled SEB reconstruction algorithm is compared to the conventional direct film coefficient reconstruction method, where it is shown that the decoupled SEB method is similarly accurate to the conventional method, but notably cheaper to evaluate. The utility of the decoupled SEB method for sensitivity analysis is demonstrated by reconstructions of a subset of the MEDLI flight data. A Monte Carlo of decoupled SEB reconstructions shows that previously-estimated uncertainties on shock layer radiation are likely inconsistent with flight data on the higher bound, even accounting for variation in surface optical properties. More detailed sensitivity analysis including much more computationally-expensive variations of in-depth properties is required to confirm this observation.

It was noted that the thermochemical solution kernel is the core of the decoupled SEB code, and future efforts in this area will be focused on improving this aspect of the code. Efforts will continue to implement a robust and efficient chemical equilibrium solver that can incorporate the ACE-style kinetic model and permit inclusion of multiple heterogeneous reactions.

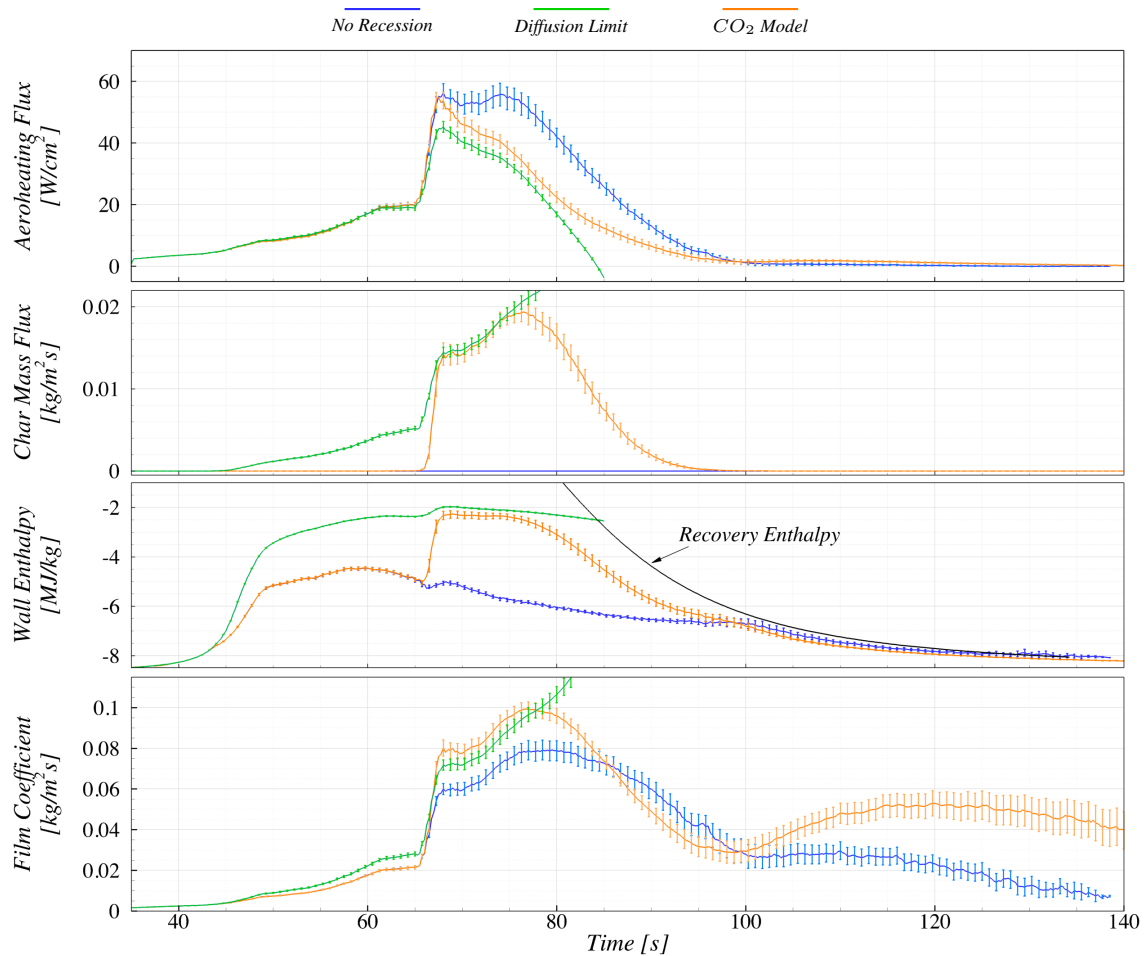


Figure 20: Reconstructed environments on MISP 7.

VII. Acknowledgements

The author would like to gratefully acknowledge that this work has relied heavily on input and material provided by Ben Blackwell, Milad Mahzari, and Todd White. Additionally, Gregory Blaisdell and Anastasios Lyrantzis have provided many helpful reviews of the material that went into this paper.

References

- ¹ D. B. Lee and W. D. Goodrich. The Aerothermodynamic Environment of the Apollo Command Module During Superorbital Entry. Technical Report NASA TN D-6792, Manned Spacecraft Center, Houston, TX 77058, 1972.
- ² F. M. Cheatwood, D. Bose, C. D. Karlgaard, C. A. Kuhl, J. A. Santos, and M. J. Wright. Mars Science Laboratory (MSL) Entry, Descent, and Landing Instrumentation (MEDLI): Complete Flight Data Set. Technical Report NASA/TM-2014-218533, NASA Langley Research Center, Langley Station, Hampton, VA, October 2014.
- ³ D. Bose, T. White, J. A. Santos, J. Feldman, M. Mahzari, M. Olson, and B. Laub. Initial Assessment of Mars Science Laboratory Heatshield Instrumentation and Flight Data. AIAA Paper 2013-0908, 51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, Grapevine, TX, January 7-10, 2013.
- ⁴ T. R. White, I. Cozmuta, J. A. Santos, B. Laub, and M. Mahzari. Proposed Analysis Process for Mars Science Laboratory Heat Shield Sensor Plug Flight Data. AIAA Paper 2011-3957, 42nd AIAA Thermophysics Conference, Honolulu, HI, June 27-30, 2011.
- ⁵ T. R. White, M. Mahzari, D. Bose, and J. A. Santos. Post-flight Analysis of the Mars Science Laboratory Entry Aerothermal Environment and Thermal Protection System Response. AIAA Paper 2013-2779, 44th AIAA

- Thermophysics Conference, San Diego, CA, June 24-27, 2013.
- ⁶ K. A. Woodbury. Sequential function specification method using future times for function estimation. In Keith A. Woodbury, editor, *Inverse Engineering Handbook*, chapter 2. CRC Press, Boca Raton, FL, 2003.
 - ⁷ A. B. Oliver and A. J. Amar. Inverse Heat Conduction Methods in the (CHAR) Code for Aerothermal Flight Data Reconstruction. AIAA Paper 2016-3384, 46th AIAA Thermophysics Conference, Washington, DC, June 13-17, 2016.
 - ⁸ A. Brandon Oliver. *Development of Inverse Methods for Reconstruction of Flight Environments on Ablators*. PhD thesis, Purdue University, December 2016.
 - ⁹ J. V. Beck, B. F. Blackwell, and C. R. St. Clair. *Inverse Heat Conduction Problems*. Wiley-Interscience, New York, NY, 1985.
 - ¹⁰ Keith A. Woodbury, editor. *Inverse Engineering Handbook*. CRC Press, Boca Raton, FL, 2003.
 - ¹¹ M. N. Ozisik and H. R. B. Orlande. *Inverse Heat Transfer: Fundamentals and Applications*. Taylor & Francis, Washington, D.C., 2000.
 - ¹² C. B. Moyer and R. A. Rindal. An Analysis of the Coupled Chemically Reacting Boundary Layer and Charring Ablator, Part II: Finite Difference Solution for the In-Depth Response of Charring Materials Considering Surface Chemical and Energy Balances. Technical Report 66-7 Part II, Aerotherm, March 1967.
 - ¹³ A. J. Amar, B. F. Blackwell, and J. R. Edwards. One-Dimensional Ablation Using a Full Newton's Method and Finite Control Volume Procedure. *AIAA Journal of Thermophysics and Heat Transfer*, 22(71-82):1, 2006.
 - ¹⁴ A. J. Amar, B. F. Blackwell, and J. R. Edwards. Development and Verification of a One-Dimensional Ablation Code Including Pyrolysis Gas Flow. *AIAA Journal of Thermophysics and Heat Transfer*, 23(59-71):1, 2006.
 - ¹⁵ A. J. Amar, N. D. Calvert, and B. S. Kirk. Development and Verification of the Charring Ablating Thermal Protection Implicit System Solver. AIAA Paper 2011-144, 49th AIAA Aerospace Sciences Meeting, Orlando, FL, January 4-7, 2011.
 - ¹⁶ A. J. Amar, A. B. Oliver, B. S. Kirk, and G. Salazar. Overview of the CHarring Ablator Response (CHAR) Code. AIAA Paper 2016-3385, 46th AIAA Thermophysics Conference, Washington, DC, June 13-17, 2016.
 - ¹⁷ Y.-K. Chen Frank Milos. Ablation and thermal response program for spacecraft heatshield design. *Journal of Spacecraft and Rockets*, 36(3):475-483, 1999.
 - ¹⁸ R. M. Kendall, R. A. Rindal, and E. P. Bartlett. Thermochemical Ablation. AIAA Paper 65-642, AIAA Thermophysics Specialist Conference, Monterey, CA, September 13-15, 1965.
 - ¹⁹ M. Mahzari, R. D. Braun, T. R. White, and D. Bose. Preliminary Analysis of the Mars Science Laboratory's Entry Aerothermodynamic Environment and Thermal Protection System Performance. AIAA Paper 2013-0185, 51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, Grapevine, TX, January 7-10, 2013.
 - ²⁰ M. Mahzari, R. D. Braun, T. R. White, and D. Bose. Inverse Estimation of the Mars Science Laboratory Entry Aerothermal Environment and Thermal Protection System Response. AIAA Paper 2013-2780, 44th AIAA Thermophysics Conference, San Diego, CA, June 24-27, 2013, 2013.
 - ²¹ B. Blackwell and J. V. Beck. A technique for uncertainty analysis for inverse heat conduction problems. *International Journal of Heat and Mass Transfer*, 53(4):753-759, January 2010.
 - ²² A. J. Amar and B. S. Kirk and A. B. Oliver and G. Salazar and J. C. Droba. Development and Verification of the CHarring Ablator Response (CHAR) Code. NASA Johnson Space Center, Houston, TX 77058, 2013.
 - ²³ M. Mahzari and R. D. Braun. Time-Dependent Mars Entry Aeroheating Estimation from Simulated In-Depth Heat Shield Temperature Measurements. *Journal of Thermophysics and Heat Transfer*, 27(3):435-446, 2013.
 - ²⁴ J. I. Frankel and M. Keyhani. Inverse Ablation Analysis and the Calibration Integral Equation Method. *Journal of Thermophysics and Heat Transfer*, Posted online on September 16, 2015.
 - ²⁵ R. M. Kendall. An Analysis of the Coupled Chemically Reacting Boundary Layer and Charring Ablator, Part V: A General Approach to the Thermochemical Solution of Mixed Equilibrium-Nonequilibrium, Homogeneous and Heterogeneous Systems. Technical Report 66-7 Part V, Aerotherm, June 1968.
 - ²⁶ F. S. Milos and Y.-K. Chen. Comprehensive Model for Multicomponent Ablation Thermochemistry. AIAA Paper 97-0141, AIAA 35th Aerospace Sciences Meeting & Exhibit, Reno, NV, January 6-9, 1997.
 - ²⁷ S. M. Scala. The Ablation of Graphite in Dissociated Air, Part I: Theory. IAS National Summer Meeting, Los Angeles, CA, June 19-22, 1962.
 - ²⁸ B. F. Blackwell and M. A. Howard. An Element Potential Based Chemical Equilibrium Solver for Gas / Surface Thermochemistry. AIAA Paper 2012-0815, 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, Nashville, TN, January 9-12, 2012.
 - ²⁹ M. E. Ewing and D. A. Isaac. Mathematical Modeling of Multiphase Chemical Equilibrium. *Journal of Thermo-*

- physics and Heat Transfer*, 29(3):551–562, 2015.
- ³⁰ J. B. Scoggins and T. E. Magin. Gibbs function continuation for linearly constrained multiphase equilibria. *Combustion and Flame*, 162(12):4514–4522, 2015.
 - ³¹ S. B. Pope. Gibbs function continuation for the stable computation of chemical equilibrium. *Combustion and Flame*, 139(3):222–226, 2004.
 - ³² J. B. Scoggins and T. E. Magin. Development of Mutation++: Multicomponent Thermodynamics And Transport properties for IONized gases library in C++. AIAA Paper 2014-2966, 11th AIAA/ASME Joint Thermophysics and Heat Transfer Conference, Atlanta, GA, June 16-20, 2014.
 - ³³ J. Lachaud, A. Martin, I. Cozmuta, and B. Laub. Ablation test-case series 1. *4th AFOSR/SNL/NASA Ablation Workshop*, 2010.
 - ³⁴ W. Kays, M. Crawford, and B. Weigand. *Convective Heat and Mass Transfer*. McGraw-Hill, 4th edition, 2005.
 - ³⁵ B. A. Cruden, A. M. Brandis, T. R. White, and M. Mahzari. Radiative Heating for MSL Entry : Verification of Simulations from Ground Test to Flight Data. AIAA Paper 2015-1894, 53rd AIAA Aerospace Sciences Meeting, Kissimmee, FL, January 5-9, 2015.
 - ³⁶ D. Bose, T. White, M. Mahzari, and K. Edquist. Reconstruction of Aerothermal Environment and Heat Shield Response of Mars Science Laboratory. *Journal of Spacecraft and Rockets*, 51(4):1174–1184, 2014.
 - ³⁷ E. Gulbransen, K. Andrew, and F. Brassart. Reaction of graphite with carbon dioxide at 1000-1600C under flow conditions. *Carbon*, 2(4):421–429, 1965.